arXiv:1410.5039v2 [math.CO] 7 Jun 2015

Cylindric Young Tableaux and their Properties

Eric Neyman
Princeton University[1]
Princeton, New Jersey, USA

Mentor: Darij Grinberg
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA

[1]Member of the Princeton undergraduate class of 2019. Most of this paper was written while the author was a student at Montgomery Blair High School in Silver Spring, Maryland.

Contents

1. Introduction **4**

2. Preliminary Definitions **4**

3. Forward Internal Insertion and Multi-Insertion **8**
3.1. Forward Insertion Algorithms and Examples 8
3.2. The Cylindric Row-Bumping Lemma and its Various Corollaries 18

4. Reverse Insertion and Reverse Multi-Insertion **22**
4.1. Reverse Insertion Algorithms and Examples 22
4.2. Relating Reverse Insertion to Forward Insertion 27
4.3. More Results about Reverse Insertion 33

5. The Cylindric RSK Correspondence **35**
5.1. The Correspondence 35
5.2. Consequences of the Cylindric RSK Correspondence 43
5.3. The Symmetry Property of CRSK 44

6. A Marble-Game Interpretation of Cylindric Tableaux **48**

7. Applying Results Concerning Cylindric Tableaux to Skew Tableaux **51**

8. A Note on Knuth Equivalence for Cylindric Tableaux **55**
8.1. Words and Knuth Equivalence 55
8.2. Cyclic Knuth Equivalence 56

Cylindric Young Tableaux and their Properties

Abstract. Cylindric Young tableaux are combinatorial objects that first appeared in the 1990s. A natural extension of the classical notion of a Young tableau, they have since been used several times, most notably by Gessel and Krattenthaler [GesKra] and by Alexander Postnikov [Post, §3]. Despite this, relatively little is known about cylindric Young tableaux. This paper is an investigation of the properties of this object. In this paper, we extend the Robinson-Schensted-Knuth correspondence, a well-known and very useful bijection concerning regular Young tableaux, to be a correspondence between pairs of cylindric tableaux. We use this correspondence to reach further results about cylindric tableaux. We then establish an interpretation of cylindric tableaux in terms of a game involving marble-passing. Next, we demonstrate a generic method to use results concerning cylindric tableaux in order to prove results about skew Young tableaux. We finish with a note on Knuth equivalence and its analog for cylindric tableaux.

Keywords: *Young tableau, cylindric tableau, skew tableau, partition, insertion, RSK Correspondence, Schur polynomial, Knuth equivalence.*

1. Introduction

In 1997, Gessel and Krattenthaler introduced cylindric semistandard Young tableaux as a modification of the concept of semistandard Young tableaux [GesKra]. They are essentially semistandard Young tableaux on a cylinder (the precise structure will be discussed in the following section). Hence, cylindric Young tableaux do not have a top row or a bottom row and, in general, are more "symmetric" than regular Young tableaux.

Being a natural extension of the notion of Young tableaux, cylindric tableaux have been studied by mathematicians during the last two decades. Among these are Alexander Postnikov [Post, §3] and Peter McNamara [McN], as well as Jennifer Morse and Anne Schilling [MorSch, §3.1]. The purpose of this paper is to expand this field through new results and applications.

In this paper, we begin by defining multi-insertion and reverse multi-insertion for cylindric tableaux, processes analogous to row-insertion and row-deletion for regular tableaux. We then prove a cylindric analog of the row-bumping lemma, a useful fact in tableau theory. We proceed to describe and prove an analog of the Robinson-Schensted-Knuth (RSK) correspondence for cylindric tableaux. This is the principal result of the paper, as the RSK correspondence for regular tableaux, which is a bijection between pairs of tableaux and matrices, has a variety of combinatorial applications [Ful, §4.3]. We adapt one of these applications — the Cauchy identity — to cylindric Schur polynomials (defined in the following section). We proceed to prove a surprising symmetry property of our cylindric RSK correspondence.

We then demonstrate an interpretation of cylindric tableaux in terms of people passing marbles in circles. We note a possible application of this interpretation to information theory. Next, we show how results concerning cylindric tableaux may be used to prove analogous results concerning regular tableaux (including skew Young tableaux).

Finally, we define a natural variation on Knuth equivalence for words that represent cylindric tableaux, which we call cyclic Knuth equivalence. We demonstrate that, despite the naturality of this extension, cyclic Knuth equivalence is not a useful construct for cylindric tableaux, as all tableaux of the same weight are cyclic Knuth equivalent under our definition. We believe that finding a variation on Knuth equivalence that does not place all cylindric tableaux of the same content in the same equivalence class would be an important development in cylindric tableau theory.

2. Preliminary Definitions

Definition 2.1. Fix positive integers k and n, with $n > k$, for the rest of this paper, excluding the examples given in the paper. The *cylinder* $\mathcal{C}_{k,n}$ is the quotient $\mathbb{Z}^2/(-k, n-k)\mathbb{Z}$. In other words, $\mathcal{C}_{k,n}$ is the set of equivalence classes of points modulo a shift by the $\mathbb{Z}$ vector $(-k, n-k)$ [Post, §3].[2]

Definition 2.2. A *cylindric partition* λ on $\mathcal{C}_{k,n}$ is a weakly decreasing sequence of integers $\ldots, \lambda_{-1}, \lambda_0, \lambda_1, \ldots$, infinite in both directions, such that for any integer m, $\lambda_m = \lambda_{m+k}+n-k$.

[2]The fact that the shift is described as $-k$ and $n-k$, instead of $-k$ and m for some m, is a standard in cylindric tableau theory [Post, §3]; the reason for this is not in the scope of this paper.

In this paper, we will think of cylindric partitions in terms of their pictures in the plane — analogues of Young diagrams. Drawn below is a cylindric partition; here, $k = 7$ and $n = 12$, and the part of the sequence of the partition that is shown below is 6, 4, 2, 2, 2, 2, 1, 1, -1, -3, -3, -3, -3, -4.

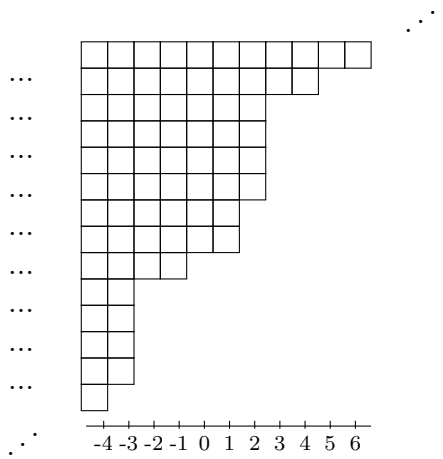

Definition 2.3. The set of cylindric partitions will be denoted Cylpar.

In this paper, the term "partition" will be used to refer to cylindric partitions, unless stated otherwise. The term "regular" will be used as the negation of "cylindric" (e.g. non-cylindric partitions will be called regular partitions).

Definition 2.4. A *point* is a pair of integers (x, y). In the diagrams in this paper, a point (x, y) is represented on the Cartesian plane by a square of side length 1. Note that the positive x-axis points downward and the positive y-axis points to the right, as usual in the theory of Young diagrams. We say that a point $P = (x, y)$ is in a partition λ (denoted $P \in \lambda$) if $y \leq \lambda_x$. Visually, P is in λ if it lies inside (i.e. to the left of the right boundary of) λ when λ is drawn on the plane as above.

Definition 2.5. A *plane row* is a set of all points with the same x-coordinate. We say that the point (x, y) is in plane row x. A *row* is the projection of a plane row onto $\mathcal{C}_{k,n}$. Thus, row x (i.e. the projection of plane row x onto the cylinder) is the same as row $x + mk$ for any integer m.

Definition 2.6. A *plane column* is a set of all points with the same y-coordinate. We say that the point (x, y) is in plane column y. A *column* is the projection of a plane column onto $\mathcal{C}_{k,n}$. Thus, column y (i.e. the projection of plane column y onto the cylinder) is the same as column $y + m(n - k)$ for any integer m.

Definition 2.7. A *box* is the projection of a point (x, y) onto $\mathcal{C}_{k,n}$. Thus, the box with coordinates (x, y) is the same as the box with coordinates $(x - mk, y + m(n - k))$, for any $m \in \mathbb{Z}$.

Definition 2.8. For any point P, $\pi(P)$ is the projection of P onto $\mathcal{C}_{k,n}$. For any box B in $\mathcal{C}_{k,n}$, $\pi^{-1}(B)$ is the set of all points P such that $\pi(P) = B$. For any plane row r, $\pi(r)$ is the projection of r onto $\mathcal{C}_{k,n}$. For any row s in $\mathcal{C}_{k,n}$, $\pi^{-1}(s)$ is the set of all plane rows r such that $\pi(r) = s$.

Definition 2.9. For any box B and plane row r such that $B \in \pi(r)$, we define $\pi_r^{-1}(B)$ to be the point in $\pi^{-1}(B)$ that is in plane row r.

Definition 2.10. A box B is in a partition λ (denoted $B \in \lambda$) if $P \in \lambda$ for some $P \in \pi^{-1}(B)$. Note that, because of the periodicity of cylindric partitions, if $P \in \lambda$ for some $P \in \pi^{-1}(B)$, then $P \in \lambda$ for all $P \in \pi^{-1}(B)$.

Definition 2.11. Let λ and μ be two cylindric partitions. We say that $\mu \subseteq \lambda$ if, for all integers m, we have $\mu_m \leq \lambda_m$.

Definition 2.12. For any cylindric partitions λ and μ such that $\mu \subseteq \lambda$, a box B is in λ/μ (denoted $B \in \lambda/\mu$) if $B \in \lambda$, but $B \notin \mu$. (Note that the set of boxes in λ/μ is finite, since there are k rows and in each row the i'th row has $\lambda_i - \mu_i$ boxes in λ/μ.) Similarly, point P is in λ/μ (also denoted $P \in \lambda/\mu$) if $P \in \lambda$, but $P \notin \mu$.

Definition 2.13. Given two cylindric partitions λ and μ such that $\mu \subseteq \lambda$, a *semistandard cylindric tableau* with *outer shape* λ and *inner shape* μ is a map R from the set of all boxes in λ/μ to a totally ordered set A such that

(a) $R(\pi((x, y_1))) \leq R(\pi((x, y_2)))$ for any x, y_1, and y_2 such that (x, y_1) and (x, y_2) are in λ/μ, and $y_1 < y_2$; and

(b) $R(\pi((x_1, y))) < R(\pi((x_2, y)))$ for any x_1, x_2, and y such that (x_1, y) and (x_2, y) are in λ/μ, and $x_1 < x_2$.

The inner and outer shapes of a semistandard cylindric tableau are regarded as part of the tableau's data; that is, the tableau "remembers" its λ and μ. A semistandard cylindric tableau can be drawn on the plane, with each square (x, y) holding the entry that the corresponding box $\pi((x, y))$ maps to under the tableau (see the diagram below). Visually, a semistandard cylindric tableau's entries increase weakly from left to right along its rows and increase strictly from top to bottom along its columns. We say that R is *bounded* by λ and μ, and that the *shape* of R is λ/μ.[3] We call A the *alphabet* of R. Frequently, alphabets of semistandard cylindric tableaux are $\mathbb{Z}^+$ or $\mathbb{Z}$. For the rest of this paper, the alphabets of all tableaux will be implicit. The elements of an alphabet are referred to as *letters*. For any particular semistandard cylindric tableau, the image of a box under the tableau is referred to as the *entry* in the box; the images of the boxes of λ/μ under the tableau are collectively called the *entries* of the tableau.

Note that, if $\mu \not\subseteq \lambda$, then there are no semistandard cylindric tableaux of shape λ/μ.

For the rest of this paper, unless stated otherwise, the word "tableau" will be used to refer to semistandard cylindric tableaux.

Below is an example of a cylindric tableau. In this example, $k = 3$ (i.e. the vertical period of the tableau is 3) and $n = 6$ (which means that the horizontal period of the tableau is $6 - 3 = 3$). The entry in a given box of a cylindric tableau in the diagram is the image of the box under the tableau map. In such diagrams, the top row that is drawn is row 0 of the tableau.

[3] While it is often convenient to think of a tableau's shape as the set of boxes between two partitions (hence the notation), λ/μ here is a notation for the pair (λ, μ); the pair of partitions may carry more information than simply the set of boxes between them.

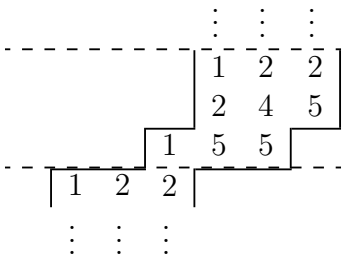

Notice that the top row that we draw is repeated at the bottom of the diagram. In this paper, we will draw tableaux as they are drawn above: the top row will be repeated at the bottom, separated by a dashed line. This convention allows us to distinguish, for example, between the following two tableaux, the first of which has $k = 1$ and the second of which has $k = 2$.

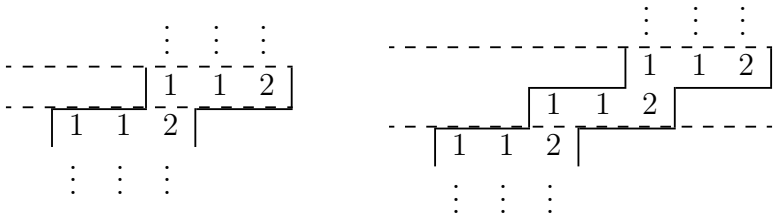

Definition 2.14. The set of semistandard cylindric tableaux of shape λ/μ, for $\lambda, \mu \in \text{Cylpar}$, will be denoted $\text{SSCT}(\lambda/\mu)$.

Definition 2.15. Given $R \in \text{SSCT}(\lambda/\mu)$, a box B is in R (denoted $B \in R$) if $B \in \lambda/\mu$.

Definition 2.16. Let R be a cylindric tableau with alphabet A. The *weight* of R (also known as the *content* of R) is the map $\text{wt}(R) : A \to \mathbb{N}$, where $\text{wt}(R)(i)$ is the size of the preimage of i under R (that is, it is the number of boxes that, in R, contain i). Here, we use $\mathbb{N}$ to represent the set of natural numbers, including 0.

When $A = \{1, 2, 3, \dots\}$, the weight of R is written as the sequence $(\text{wt}(R)(1), \text{wt}(R)(2), \dots)$, and can be truncated at a place where all following terms are zero. For example, the cylindric tableau shown below has weight $(1, 3, 1, 0, 1)$.

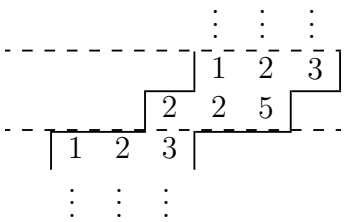

Definition 2.17. Given an alphabet A, for every $a \in A$, let x_a be a variable. For any two $a, b \in A$, x_a and x_b are distinct variables. We will denote the family $(x_a)_{a \in A}$ by $\mathbf{x}$; it will be called a *variable set*.

Definition 2.18. Given a tableau R, we will define the *weight monomial* of R with variable set $\mathbf{x}$, denoted $\mathbf{x}^{\text{wt}(R)}$, as follows:

$$\mathbf{x}^{\text{wt}(R)} = \prod_{a \in A} {x_a}^{\text{wt}(R)(a)} = \prod_{B, \text{ box in } R} x_{R(B)}.$$

We use $R(B)$ above to denote the entry of R in B.

For example, if R is the tableau that is shown as an example in Definition 2.16, then $\mathbf{x}^{\mathrm{wt}(R)} = x_1 x_2^3 x_3 x_5$.

Definition 2.19. Given two cylindric partitions μ and λ such that $\mu \subseteq \lambda$, the *Schur polynomial* of λ/μ with variable set $\mathbf{x}$, denoted $s_{\lambda/\mu}(\mathbf{x})$, is a power series of bounded degree, defined as follows:

$$s_{\lambda/\mu}(\mathbf{x}) = \sum_{R \in \mathrm{SSCT}(\lambda/\mu)} \mathbf{x}^{\mathrm{wt}(R)}.$$

Definition 2.20. A *horizontal strip* is a pair λ/μ of partitions μ and λ, with $\mu \subseteq \lambda$, such that none of the boxes in λ/μ are in the same column. We say that a set S of boxes forms a horizontal strip if there exists a horizontal strip h such that S is the set of boxes in h.[4]

3. Forward Internal Insertion and Multi-Insertion

3.1. Forward Insertion Algorithms and Examples

Definition 3.1. Given $R \in \mathrm{SSCT}(\lambda/\mu)$, a box $B = \pi((i,j))$ is an *inside cocorner* of R if $B \notin \mu$, but $\pi((i-1,j)) \in \mu$ and $\pi((i,j-1)) \in \mu$. Visually, B is an inside cocorner of R if it is to the right of (outside) μ, but the boxes to the left of and above B are inside μ. Note that B is not necessarily a box of R, as it is possible that B is outside λ as well.

Given a cylindric tableau R and an inside cocorner B of R, we can *internally insert* B into R, in a process that is similar to regular row-insertion, using the algorithm below, which takes a cylindric tableau and an inside cocorner of the tableau as input and outputs a new tableau (we will later prove that the output is indeed a valid semistandard tableau). This algorithm is inspired by the internal row insertion for skew tableaux described by Bruce Sagan and Richard Stanley [SagStan, §2].

Algorithm 3.2 (Internal Row-Insertion)**.**

Function Insert(tableau R, box B) ▷ B must be an inside cocorner of R.
1: $\mu :=$ inner shape of R.
2: $\lambda :=$ outer shape of R.
3: **if** $B \in R$ **then**:
4: $x :=$ entry of R that is in B.
5: **end if**.
6: Expand μ to include B and remove B from R.
7: **if** $B \notin \lambda$ **then**: ▷ This happens only when B was not in R to begin with.
8: Expand λ to include B.
9: **else**:
10: **while** $x \neq$ null **do**:
11: $r :=$ row of B.

[4]It is worth noting that λ/μ is a horizontal strip if and only if $\lambda_i \geq \mu_i \geq \lambda_{i+1}$.

12: **if** x is greater than or equal to every entry in R in row $r+1$ **then**:
13: $B :=$ leftmost box of row $r+1$ that is not in λ.
14: Put x in B and expand λ to include B. ▷ We say that x *lands* in B.
15: $x :=$ null.
16: **else**:
17: $B :=$ box of the leftmost entry of R in row $r+1$ that is greater than x.
18: $x' :=$ entry of R in B.
19: In R, replace x' in B with x.
20: $x := x'$.
21: **end if**.
22: **end while**.
23: **end if**.
24: **return** R.

Note that, although we use the phrase, "internally insert B into R," the result of the process is that B, which was previously in R, is no longer in R.

Example 3.3. Suppose we want to apply Algorithm 3.2 to (R, B), where R is the tableau shown below and B is the box that contains 1 in R. We say that we *bump* the 1 from B (or that the 1 is *bumped* from B).

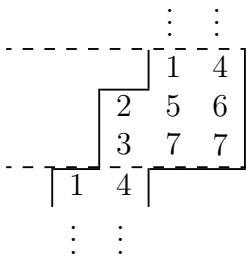

We expand μ to include the box containing the 1 and exclude the 1 from the tableau.

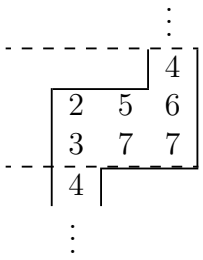

We then replace the leftmost entry greater than 1 in the following row (which is 2) with 1. We say that the 2 is *bumped* by the 1.

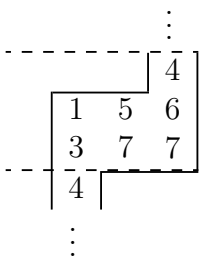

This process continues as follows:

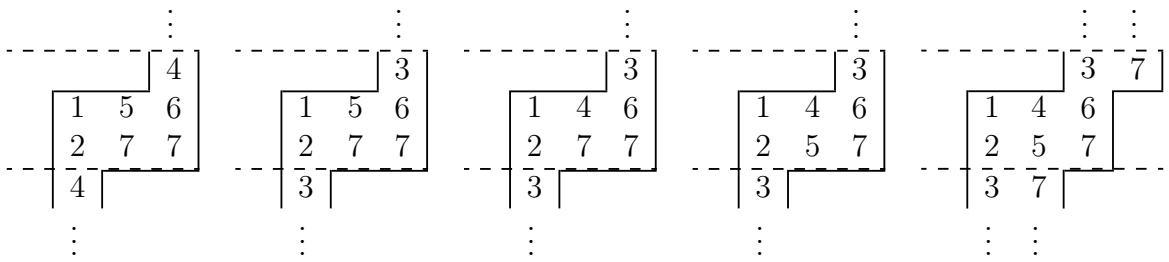

In the final step, there is no entry greater than 7 in the following row, and the 7 lands at the end of the following row, to the right of the 3. The insertion process is now complete.

Remark 3.4. *Consider any box C of R. Since on line 19 of Algorithm 3.2 an entry only replaces an entry greater than it, it follows that the entry in box C can only decrease (or disappear) throughout the row-insertion process, and does decrease if and when an entry is bumped out of C and a new entry takes its place.*

We now draw the final tableau above, this time with more repeating rows in the diagram.

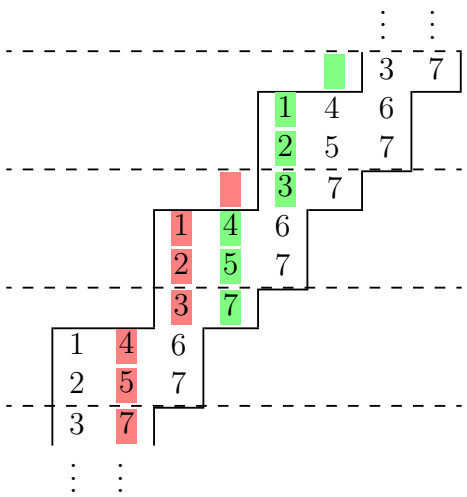

Above in green and red are two bumping routes, which we will define shortly.

Remark 3.5. *Algorithm 3.2 always ends after a finite number of steps.*

Proof. Let M be the largest entry of R. Every iteration of the loop beginning on line 10, x increases; however, x cannot be greater than M. Thus, the loop must terminate and the algorithm must end. □

Definition 3.6. Let R be a tableau and P be a point such that $\text{Insert}(R, \pi(P))$ is well-defined. We will define the *bumping route* of P as a list of points (not boxes — and hence the two bumping routes above are distinct) that is constructed as follows:

Suppose that we perform Algorithm 3.2 with R and $\pi(P)$ as input. We will modify the algorithm for the purposes of this definition. At the very beginning, we let s be the plane

row of P, and we initialize H as a list with the single entry $P = \pi_s^{-1}(\pi(B))$ (see Definition 2.9). At the very end of the loop starting on line 10, add a line that increments s by 1.[5] Immediately after the new line, append $\pi_s^{-1}(B)$ to the end of H. The bumping route of P is defined to be the list H as it stands after the algorithm terminates.

Intuitively, then, the bumping route is a list of points in consecutively increasing rows, such that each point corresponds to a value of B in Algorithm 3.2 during its execution.

Later, we will prove two important results: that Algorithm 3.2 necessarily outputs a valid semistandard tableau and that each box in a bumping route (except the first one) is weakly left of the box before it (we say that bumping routes *trend weakly left*).

It is also possible to internally row-insert multiple entries at the same time. The algorithm below is the algorithm for one-step multi-insertion. It takes a tableau and a regular insertion queue (defined below).

Definition 3.7. A *queue* is a data structure that operates on a "first-in-first-out" basis: when elements are added to a queue, they are added to the end of the queue, but when elements are removed from a queue, they are removed from the beginning of the queue. An *insertion queue* is a queue of pairs, such that the first element of each pair is a letter (element of the (implicit) alphabet) and the second element of each pair is a row.

Definition 3.8. An insertion queue is *regular* if, for any two elements of the queue (x_1, r) and (x_2, r), where $x_1 < x_2$, (x_1, r) comes before (x_2, r) in the queue.[6]

The algorithm below (the subroutine to our multi-insertion algorithm) outputs a pair, whose first element is a map from a subset of $\mathcal{C}_{k,n}$ to the alphabet of the tableau taken as input (though the map itself is not necessarily a valid semistandard tableau), and whose second element is an insertion queue corresponding to the entries bumped from the tableau via insertion of the entries in the insertion queue that is taken as a parameter.

Algorithm 3.9 (One-Step Multi-Insertion)**.**

Function OneStepMulti(tableau R, insertion queue $\mathfrak{q}$) ▷ $\mathfrak{q}$ must be regular.

1: $\lambda :=$ outer shape of R.
2: $\mathfrak{q}' :=$ empty insertion queue.
3: **while** $\mathfrak{q}$ is not empty **do**:
4: Remove the first element from $\mathfrak{q}$. Let it be (x, r).
5: **if** x is greater than or equal to every entry in R in row r **then**:
6: $B :=$ leftmost box of row r that is not in λ.
7: Put x in B and expand λ to include B. ▷ λ is not necessarily a valid partition anymore.
8: **else**:
9: $B :=$ box of the leftmost entry of R in row r that is greater than x.
10: $x' :=$ entry of R in B.

[5] After this augmentation, $B \in \pi(s)$.

[6] The fact that the second element of a pair in an insertion queue is a row should be kept in mind. For example, if $k = 3$, $(3, 5)$ cannot come before $(2, 2)$ in a regular insertion queue, since 2 and 5 would refer to the same row.

11: In R, replace x' in B with x.
12: Add $(x', r+1)$ to $\mathfrak{q}'$.
13: **end if**.
14: **end while**.
15: **return** $(R, \mathfrak{q}')$. ▷ R is not necessarily a tableau.

Remark 3.10. *The insertion queue returned by Algorithm 3.9 is regular.*

Proof. Let $(R', \mathfrak{q}') = \text{OneStepMulti}(R, \mathfrak{q})$ for a tableau R and a regular insertion queue $\mathfrak{q}$. Suppose, for contradiction, that $\mathfrak{q}'$ is not regular. Then there exist two elements of $\mathfrak{q}'$, (y_1, r) and (y_2, r), such that $y_1 < y_2$, but (y_2, r) comes before (y_1, r) in $\mathfrak{q}'$. Let x_1 and x_2 be the entries that bumped out y_1 and y_2, respectively. Then $(x_2, r-1)$ and $(x_1, r-1)$ were in $\mathfrak{q}$, with $(x_2, r-1)$ coming first. Since $\mathfrak{q}$ is regular, it follows that $x_2 < x_1$. We also know that $x_1 < y_1$, since x_1 bumps out y_1. It follows that $x_2 < x_1 < y_1 < y_2$. Since $y_1 < y_2$, y_1 is to the left of y_2 in R, and $y_1 > x_2$. It follows that y_2 is not the leftmost entry of R in its row that is greater than x_2 at the time that x_2 is to be inserted, which is a contradiction, because then x_2 does not bump out y_2. Thus, $\mathfrak{q}'$ is a regular insertion queue. □

Remark 3.11. *Let R be a cylindric tableau and let $\mathfrak{q}_1$ and $\mathfrak{q}_2$ be two regular insertion queues that are permutations of each other. Let $(R_1, \mathfrak{q}_1') = \text{OneStepMulti}(R, \mathfrak{q}_1)$ and $(R_2, \mathfrak{q}_2') = \text{OneStepMulti}(R, \mathfrak{q}_2)$. Then $R_1 = R_2$ and $\mathfrak{q}_1'$ is a permutation of $\mathfrak{q}_2'$.*

Proof. Consider a particular row r and all of the pairs of $\mathfrak{q}_1$ and $\mathfrak{q}_2$ to be placed in r (in other words, all pairs whose second element is r). These are the same pairs, because $\mathfrak{q}_1$ is a permutation of $\mathfrak{q}_2$. Furthermore, regularity defines an ordering (least to greatest) among these pairs based on their first elements, so these pairs are also in the same order. Since insertion into row r is only affected by the entry being inserted into row r and the entries already in row r, it follows that row r is the same in R_1 and R_2 after the two insertion queues have been processed. Since this is true for all r, it follows that $R_1 = R_2$.

Since $R_1 = R_2$, the same entries were bumped out of R to produce R_1 and R_2 (though not necessarily in the same order). These bumped-out entries, along with the row numbers of the rows immediately below the ones from which they were bumped out, constitute the pairs in $\mathfrak{q}_1'$ and $\mathfrak{q}_2'$, respectively. These pairs being the same, it follows that $\mathfrak{q}_1'$ is a permutation of $\mathfrak{q}_2'$. □

Finally, using Algorithm 3.9, we can internally row-insert multiple boxes into a tableau simultaneously. The following algorithm does this, taking a tableau and a set of boxes that forms a horizontal strip as input (with the precondition that, when the inner shape of the tableau is expanded to include these boxes, it remains a valid partition) and outputs a tableau (we will prove later that the output is indeed a valid semistandard tableau).

Algorithm 3.12 (Full Multi-Insertion)**.**

Function FullMulti(tableau R, set S of boxes) ▷ The inner shape of R plus the boxes in S must be a valid partition; no box in S is in the inner shape of R. Also, S must form a horizontal strip.

$\mu :=$ inner shape of R.
$\lambda :=$ outer shape of R.
$\mathfrak{q}_0 :=$ empty insertion queue.
Choose any integer r. ▷ We prove later that the choice of r is immaterial.
$h := r$.
while $h \neq r + k$ **do**: ▷ k is the vertical period of R.
    $L :=$ list of boxes in row r in S, from left to right.
    while L is not empty **do**:
        $B :=$ first element of L.
        Remove B from L.
        if $B \in R$ **then**:
            $x :=$ entry of R in B.
            Put $(x, h + 1)$ into $\mathfrak{q}_0$.
            Expand μ to include B and remove B from R.
        else:
            Expand μ and λ to include B and remove B from R.
        end if.
    end while.
    $h := h + 1$.
end while.
$R_0 := R$.
$i := 0$.
while $\mathfrak{q}_i$ is not empty **do**:
    $(R_{i+1}, \mathfrak{q}_{i+1}) :=$ OneStepMulti $(R_i, \mathfrak{q}_i)$. ▷ This line inserts all elements of $\mathfrak{q}_i$ into R_i and denotes the resulting tableau and queue as R_{i+1} and $\mathfrak{q}_{i+1}$, respectively.
    $i := i + 1$.
end while.
$R := R_i$.
return R.

Example 3.13. Let R, drawn below, be the input tableau into Algorithm 3.12, and let the red boxes in the diagram below constitute S.

$R =$

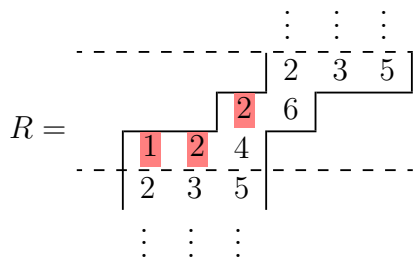

Suppose that we pick row 0 (the top row) to be our starting row r. We expand μ to include the red boxes, deleting those boxes from R and putting corresponding entries into $\mathfrak{q}_0$. The 2 is deleted from row 1, and it is to be inserted into the following row, so we add $(2, 2)$ to $\mathfrak{q}_0$. We then add $(1, 0)$ (the same as $(1, 3)$, since the vertical period of R is 3) and $(2, 0)$ into $\mathfrak{q}_0$. We let R_0 be our current tableau.

$\mathfrak{q}_0 = (2,2),(1,0),(2,0)$ $\qquad R_0 =$

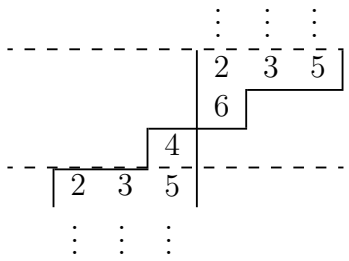

We now enter the One-Step Multi-Insertion subroutine. For this subroutine, set $\mathfrak{q} = \mathfrak{q}_0$, and let $\mathfrak{q}'$ be the empty queue. We remove $(2,2)$ from $\mathfrak{q}$ and insert 2 into row 2, bumping out the 4 and adding $(4,0)$ to $\mathfrak{q}'$. We remove $(1,0)$ from $\mathfrak{q}$ and insert 1 into row 0, bumping out the 2 and adding $(2,1)$ to $\mathfrak{q}'$. We remove $(2,0)$ from $\mathfrak{q}$ and insert 2 into row 0, bumping out the 3 and adding $(3,1)$ to $\mathfrak{q}'$. Finally, with $\mathfrak{q}$ empty, we exit the subroutine. We let $\mathfrak{q}_1$ and R_1 be the returned insertion queue and the returned tableau, respectively; they are shown below:

$\mathfrak{q}_1 = (4,0),(2,1),(3,1)$ $\qquad R_1 =$

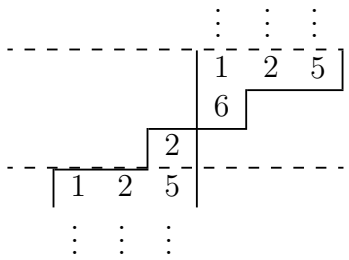

After another run through the subroutine, we have:

$\mathfrak{q}_2 = (5,1),(6,2)$ $\qquad R_2 =$

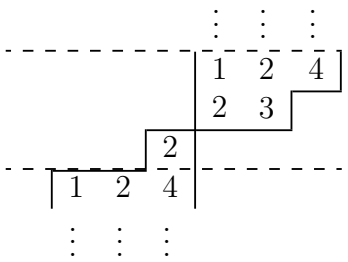

Notice that a box got added to R_2. This happened because, when $(3,1)$ was removed from $\mathfrak{q}$ in the subroutine, the largest entry of row 1 was 2, so 3 was added at the end of row 1 and nothing was bumped out. For the same reason, $\mathfrak{q}_2$ now has only two elements.

After a third iteration of the subroutine, we have:

$\mathfrak{q}_3$ is empty $\qquad R_3 =$

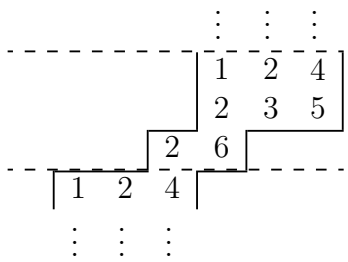

Now that $\mathfrak{q}_3$ is empty, we exit out of the loop, let R equal R_3, and return R.

If S consists of only one box, then Algorithm 3.12 functions as Algorithm 3.2. This means that when we prove certain properties of full multi-insertion (such as the fact that it returns a valid semistandard tableau), we will be showing the analogous properties to be true of single insertion as well.

Remark 3.14. *Algorithm 3.12 always ends after a finite number of steps.*

Proof. Consider any entry that is originally bumped out of R. This entry will bump out a larger entry, and that entry will bump out a larger entry, and so on. However, the bumped-out entry can only be as large as the largest entry that was originally in R. Therefore, any chain of bumps will eventually end with a landing; when this happens to every chain, the queue becomes empty and Algorithm 3.12 terminates. □

The concept of a bumping route can be extended to multi-insertion. The intuitive idea is that we keep track of an entry's element in the queue and, when we insert the entry, we put the box's corresponding point (in the following row) into the bumping route. Thus, when we perform multi-insertion on a tableau and a set of boxes, we can construct the bumping route of any point corresponding to any box in the set. Our formal extension of the definition of a bumping route to multi-insertion is below.

Definition 3.15. Let R be a tableau and S be a set of boxes such that FullMulti(R, S) is well-defined. Let P any point such that $\pi(P) \in S$. We will define the *bumping route* of P with input box set S as a list of points that is constructed as follows:[7]

Suppose that we perform Algorithm 3.12 with R and S as input. We will modify Algorithms 3.9 and 3.12 for the purposes of this definition. At the very beginning, we let s be the plane row of P, and we initialize H as a list with the single entry P. On line 13 of Algorithm 3.12, if $B = \pi^{-1}(P)$, instead of putting $(x, h+1)$ into $\mathfrak{q}_0$, put $(x, h+1, s+1)$ into the queue.[8]

In Algorithm 3.9, if in an iteration of the loop beginning on line 3 a three-element member of $\mathfrak{q}$ is removed (let it be (x, r, s)), on line 12 add $(x', r+1, s+1)$ into $\mathfrak{q}'$ instead of $(x', r+1)$. Furthermore, if a three-element member (x, r, s) is removed from the queue, at the very end of that iteration of the loop, append $\pi_s^{-1}(B)$ to the end of H.

The bumping route of P is defined to be the list H as it stands after Algorithm 3.12 terminates.

Proposition 3.16. *For any R and S satisfying the preconditions of Algorithm 3.12, after every bump that occurs during the execution of Algorithm 3.12 with inputs R and S, R is a valid semistandard tableau. Also, for every bumping route created through the insertions caused by Algorithm 3.12, each point of the bumping route (except the first) is weakly left of the point before it — that is, all thus generated bumping routes trend weakly left.*

Proof. Let A be the alphabet of R. Define QuA $= A \cup \{-\infty, \infty\}$. We compare $-\infty$ and ∞ to all elements of QuA as such: $-\infty < a < \infty$ for all $a \in$ QuA (including $a = -\infty$ and $a = \infty$). That is, we have a quasiorder on QuA: in addition to the ordering of A, we have $-\infty < a < \infty$ for all $a \in A$, as well as $-\infty < -\infty$, $-\infty < \infty$, and $\infty < \infty$.

[7]The input box set S will be implicit in all future discussion of bumping routes.

[8]For the purposes of this definition only, we extend the definition of an insertion queue to include triples of the form (letter, row, plane row).

In this proof, we will think of a tableau as a map from $\mathcal{C}_{k,n}$ (not just a part of $\mathcal{C}_{k,n}$) to QuA, where the images of all boxes in the inner shape of the tableau are $-\infty$ and the images of all boxes not in the outer shape of the tableau are ∞. A similar technique was used by Donald Knuth in his studies of regular tableaux [Knu, §3]. If R is a tableau by our regular definition, then we will denote $\widehat{R}$ to be the corresponding tableau by our new definition. Note that, by looking at $\widehat{R}$, we know R and the two partitions that bound R.

For this proof, we will modify line 9 of Algorithm 3.9 as such: "$B :=$ box of the leftmost entry of $\widehat{R}$ in row r that is *not less-than-or-equal-to* x" (note that "greater than" and "less than or equal to" are not opposites in our quasiorder). This will allow us to think of an element of S being bumped out as $-\infty$ being inserted into the corresponding row, and it will allow us to think of a box landing in a square as bumping out ∞ and the ∞ being inserted "infinitely far to the right" in the following row.

Definition 3.17. Given a tableau $\widehat{R}$ by our new definition, define the *inner shape* of $\widehat{R}$ to be the partition containing exactly the $-\infty$'s in $\widehat{R}$. Define the *outer shape* of $\widehat{R}$ to be the partition that does not contain exactly the ∞'s in $\widehat{R}$.

Our definition above is consistent with the partitions of R during the execution of Algorithm 3.12; by looking at the placement of $-\infty$ and ∞ in $\widehat{R}$, we can determine R's inner and outer shape. This consistency includes line 16 of Algorithm 3.12: $-\infty$ is placed into B, and thus B becomes part of both the inner and outer shapes of both $\widehat{R}$ and R.

Suppose that we know that, at some point during the execution of Algorithm 3.12, the entries of $\widehat{R}$ increase weakly from left to right and strictly from top to bottom (that is, $\widehat{R}$ is semistandard). Then it follows that R is bounded by two partitions (otherwise, the columns would not be strictly increasing, or the rows would fail to have the form (infinitely many $-\infty$'s, some elements of A, infinitely many ∞'s)). Thus, for the first part of our proposition, it suffices to show that, at every point during the execution of Algorithm 3.12, the entries of $\widehat{R}$ increase weakly from left to right and strictly from top to bottom.

We will prove our proposition by induction. We know that R (and therefore $\widehat{R}$) is a valid semistandard tableau to start with and that all bumping routes (all with 0 boxes thus far) trend weakly left. Suppose that R (and $\widehat{R}$) is a valid semistandard tableau after every bump up to but not including the insertion of an entry x into a box B (in the case that $B \in S$, we have $x = -\infty$ for the purposes of insertion into $\widehat{R}$). Suppose further that all bumping routes created up to this point in the execution of Algorithm 3.12 have trended weakly left. We will show that, after the insertion of x into B, R remains a valid semistandard tableau, and the bumping route that B becomes a part of after this insertion still trends weakly left.

Just before the insertion of x into B, R is semistandard. Let r be the row containing box B. Since x is put into the leftmost box containing an entry not less-than-or-equal-to x, all boxes directly to the left of B (that is, all boxes in the row of B to the left of B) have entries that are less than or equal to x; furthermore, since x is less than the entry previously in B, we have that x is less than or equal to all entries directly to the right of B. Similarly, we have that x is less than all entries directly below B (in the column of B below B). To show that R continues to be a valid semistandard tableau after the insertion of x into B, it thus remains only to show that x is greater than all entries directly above B when it is inserted into B.

Suppose that x is inserted into B during the formation of $\widehat{R_j}$, for some $j \geq 0$ (that is, on the j'th call of Algorithm 3.9 or, if $j = 0$, before line 21 of Algorithm 3.12).

Case 1: $j = 0$**.** In this case, the bumping route containing B only has B, so clearly it trends weakly left. Also, by the horizontal strip restriction on S, we have that all entries above B are $-\infty$, as desired.

Case 2: $j > 0$**.** In this case, let B' be the box that is immediately above B (in row $r-1$), let C' be the box in row $r-1$ from which the x that is inserted into B was bumped (during the formation of $\widehat{R_{j-1}}$), and let C be the box immediately below C' (in row r). First we show that the bumping route that includes B after the insertion of x into B trends weakly left after B is added to it — that is, that B is weakly to the left of C. To show this, it suffices to show that the entry in C in $\widehat{R_{j-1}}$ is not less-than-or-equal-to x (because of how B is chosen). (Note that the entry in C in $\widehat{R_{j-1}}$ is the entry in C just before x is inserted into B. This is because during the formation of $\widehat{R_j}$, pairs of the insertion queue with second element r are processed based on the first elements' boxes in row $r-1$ from left to right; since x was in C', all entries thus far processed came from boxes in $r-1$ to the left of C', and thus landed in boxes in r to the left of C, by our inductive hypothesis about bumping routes.)

Suppose, for contradiction, that the entry in C in $\widehat{R_{j-1}}$ — call it y — is less than or equal to x. Then $\widehat{R_{j-2}}$ (or $\widehat{R}$, in the case of $j = 1$), which contains x in C', cannot contain y in C. Since y is in C in R_{j-1}, but is not in C in R_{j-2}, we have that y is inserted into C during the formation of $\widehat{R_{j-1}}$. We know that $C \notin S$, because C is below C' (and S is a horizontal strip). This means that y is bumped out of row $r-1$. In fact, since y is inserted into row r during the formation of R_{j-1}, it follows that y is bumped out of row $r-1$ during the formation of $\widehat{R_{j-2}}$ — a nonsensical concept and, therefore, a contradiction if $j = 1$. If $j > 1$, then by our inductive hypothesis, we know that y is bumped from a box that is weakly to the right of C' during the formation of $\widehat{R_{j-2}}$. Hence, a box weakly to the right of C' contains an entry that is not greater-than-or-equal-to y in $\widehat{R_{j-2}}$, while C' contains x in $\widehat{R_{j-2}}$ — a contradiction, since $y \leq x$. Therefore, our inductive step holds for the part of our proposition that concerns bumping routes.

We now know that B is weakly to the left of C and that B' is weakly to the left of C'. Thus, the entry in B' right before the insertion of x into B is less than or equal to the entry in C' at this time, which is in turn less than x (since x was previously bumped out of C'). Therefore, x is greater than the entry in B' (and all the entries directly above B') when it is inserted into B, as desired.

Having completed our induction, we have proven our proposition. □

Suppose that S contains only one element (call it B). In this case, the precondition of the inner shape of R plus B being a valid partition is equivalent to B being an inside cocorner. Thus, with only one element in S, Algorithm 3.12 is reduced to Algorithm 3.2. Therefore, Algorithm 3.2 necessarily produces a valid semistandard tableau, and the bumping route of any $P \in \pi^{-1}(B)$ trends weakly left.

Remark 3.18. *The output of Algorithm 3.12 does not depend on the choice of r in line 4 of Algorithm 3.12.*

Proof. Take two values of r — call them r_1 and r_2 — and perform Algorithm 3.12 on a tableau R and a set S of boxes using r_1 and r_2. For any i, let $R_i(r_1)$ be R_i produced when the algorithm is run using $r = r_1$ and let $R_i(r_2)$ be R_i produced when the algorithm is run using $r = r_2$. Similarly, let $\mathfrak{q}_i(r_1)$ be $\mathfrak{q}_i$ produced when the algorithm is run using $r = r_1$ and let $\mathfrak{q}_i(r_2)$ be $\mathfrak{q}_i$ produced when the algorithm is run using $r = r_2$.

We will proceed by induction on i. We know that $R_0(r_1) = R_0(r_2)$ since the same boxes are taken out of R either way. We also know that $\mathfrak{q}_0(r_1)$ contains the same elements as $\mathfrak{q}_0(r_2)$, and that both insertion queues are regular, as entries were taken out from R within any particular row from left to right (and thus from least to greatest). Suppose that we have $R_g(r_1) = R_g(r_2)$ and the queues $\mathfrak{q}_g(r_1)$ and $\mathfrak{q}_g(r_2)$ are regular and are permutations of each other (contain the same elements) for all $g < i$ for some positive i. Then, by Remark 3.11, we have $R_i(r_1) = R_i(r_2)$, and the queues $\mathfrak{q}_i(r_1)$ and $\mathfrak{q}_i(r_2)$ are permutations of each other. Furthermore, both $\mathfrak{q}_i(r_1)$ and $\mathfrak{q}_i(r_2)$ are regular, by Remark 3.10. Having completed our induction, we have shown that $R_i(r_1) = R_i(r_2)$ for all i, including the final value that i takes on. Thus, the output of Algorithm 3.12 does not depend on the choice of r. □

3.2. The Cylindric Row-Bumping Lemma and its Various Corollaries

Lemma 3.19. *Let R be a tableau on which multi-insertion is performed. Then for all rows r of R, the list of entries bumped out of r in the order in which they were bumped out is weakly increasing.*

Proof. It suffices to show that for any two entries bumped consecutively from a given row s ("consecutively" meaning that nothing is bumped out from s in the meantime), the first of the two entries that is bumped is less than or equal to the second entry that is bumped. We will proceed by induction on the moment of time the second entry is bumped. Suppose that, for every row s, any two entries e_1 and e_2 that are bumped consecutively from s (in that order) satisfy $e_1 \leq e_2$, up to but not necessarily including the bumping of entries a_1 and a_2 (in that order) from a row r. We will show that $a_1 \leq a_2$.

Suppose that $a_1 \in S$ (meaning that the box containing a_1 lies in S) and $a_2 \in S$. Since a_1 is bumped first, a_1 is to the left of a_2, and so $a_1 \leq a_2$. Suppose that $a_1 \in S$ and $a_2 \notin S$. Then a_1 and a_2 are both originally in row r in R (a_2 is originally in row r because it is bumped out consecutively after a_1 is), with a_1 to the left of a_2. Again, it follows that $a_1 \leq a_2$. Clearly, it cannot be the case that $a_1 \notin S$ and $a_2 \in S$. (Note that this is why we do not need a base case for our induction: we do not use our inductive hypothesis for the entries originally bumped from R.)

Now, suppose that $a_1 \notin S$ and $a_2 \notin S$. Let b_1 and b_2 be the entries that bump out a_1 and a_2, respectively. We have that b_1 is inserted into r before b_2 is. It follows that b_1 is bumped from row $r-1$ before b_2 is (because queues work on a first-in-first-out basis). It follows by our inductive hypothesis that $b_1 \leq b_2$. Since b_1 and b_2 are inserted consecutively into row r, it follows that b_2 is placed strictly to the right of b_1. Therefore, $a_1 \leq a_2$, as desired. □

Definition 3.20. Let H be a bumping route (created by an insertion algorithm, such as Algorithm 3.2 and Algorithm 3.12). For all plane rows r such that H has a point in r, $H(r)$ will denote that point. $L(H)$ will denote the number of points in H, and for $1 \leq i \leq L(H)$, H_i will denote the i'th point of H.

Define a total ordering on points as follows: $P = (x_1, y_1) \leq Q = (x_2, y_2)$ if and only if (a) $x_1 < x_2$ (P is above Q), or (b) $x_1 = x_2$ and $y_1 \geq y_2$ (P and Q are in the same row, and P is weakly right of Q). (The $\geq$ sign is intended.)

Theorem 3.21 (Cylindric Row-Bumping Lemma)**.** *Let R be a tableau and S be a set of boxes such that R and S satisfy the input preconditions for Algorithm 3.12. Let G and H be two bumping routes created when Algorithm 3.12 is performed with R and S as input, such that $G_1 < H_1$. Then for all plane rows r such that $G(r)$ and $H(r)$ are defined, $H(r)$ is strictly to the left of $G(r)$ and the bump that extended G to include $G(r)$ came after the bump that extended H to include $H(r)$.*

Proof. Let s be the plane row that H_1 is in. Since $G_1 < H_1 = H(s)$, G_1 is in s or is above s. If $G_{L(G)}$ is in a plane row above s, then all elements of G are above all elements of H, meaning that $G(r)$ and $H(r)$ are not both defined for any r; in this case, we are done. Otherwise, we know that $G(s)$ is to the right of $H(s)$: if $G_1 = G(s)$, then G_1 must be to the right of H_1 in order for $G_1 < H_1$ to hold; if $G(s) = G_m$ for some $m > 1$, then $G(s)$ must be to the right of $H(s)$ because $\pi(H(s))$ is a box that was originally removed from the tableau, and thus $H(s)$ is permanently to the left of the tableau. Similarly, we know that the bump that extended G to include $G(s)$ came after the bump that extended H to include $H(s)$: if $G_1 = G(s)$, then $G(s)$ is to the right of $H(s)$, which means that it is bumped out later (since the original bumps proceed left to right across rows); if $G(s) = G_m$ for some $m > 1$, then G reaches plane row s on a later call of Algorithm 3.9.

Suppose that $H(r)$ is to the left of $G(r)$ and that the bump that extended G to include $G(r)$ came after the bump that extended H to include $H(r)$ for all r such that $s \leq r < t$. We are to show that $H(t)$ is to the left of $G(t)$ and that the bump that extended G to include $G(t)$ came after the bump that extended H to include $H(t)$.

Let $u = \pi(t)$. Let g and h be the entries bumped out of $\pi(G(t-1))$ and $\pi(H(t-1))$ when G and H were extended to include $G(t-1)$ and $H(t-1)$, respectively. If $G(s) = G_m$ for some $m > 1$, then G was extended to include $G(t)$ on a later call of Algorithm 3.9 than the call on which H was extended to include $H(t)$; thus, in this case, the second part of our claim is obviously true. If $G(s) = G_1$, then (h, u) and (g, u) are in the same queue, but (h, u) comes first by our inductive hypothesis; it follows that the insertion of h into row u came before the insertion of g into row u. Thus, the bump that extended G to include $G(t)$ did indeed come after the bump that extended H to include $H(t)$.

By our inductive hypothesis, h was bumped out of row $u-1$ before g was. It follows by Lemma 3.19 that $h \leq g$. Let h' be the entry in $\pi(H(t))$ at the time that g is inserted into row u. We know that $h' \leq h$ (Remark 3.4). Thus, $h' \leq g$. It follows that, at this time, g is greater than or equal to the entry in $\pi(H(t))$ and all entries to the left of $\pi(H(t))$ in row u. Thus, g is inserted in row u strictly to the right of $H(t)$. Therefore, $H(t)$ is to the left of $G(t)$, as desired.

Having completed our induction, we have shown that $H(r)$ is to the left of $G(r)$, and that the bump that extended G to include $G(r)$ came after the bump that extended H to include $H(r)$, for all plane rows r such that $G(r)$ and $H(r)$ are defined. □

Corollary 3.22. *Let G and H be two bumping routes such that, for some plane row r, $H(r)$ is strictly to the left of $G(r)$. Then $G_1 < H_1$, and for all plane rows s such that $G(s)$ and $H(s)$ are defined, $H(s)$ is strictly to the left of $G(s)$.*

Proof. Suppose, for contradiction, that $G_1 \geq H_1$. Clearly, $G_1 \neq H_1$, because otherwise G and H would be the same bumping route and it would not be the case that $H(r)$ is strictly to the left of $G(r)$ for some r. Thus, $G_1 > H_1$, meaning that $H_1 < G_1$. By Theorem 3.21, this implies that $G(r)$ is to the left of $H(r)$, a contradiction. Thus, $G_1 < H_1$. The rest of the corollary follows directly from Theorem 3.21. □

Corollary 3.23. *Let G and H be two bumping routes such that, for some plane row r, the bump that extends G to include $G(r)$ came after the bump that extended H to include $H(r)$. Then $G_1 < H_1$, and for all plane rows s such that $G(s)$ and $H(s)$ are defined, the bump that extends G to include $G(s)$ came after the bump that extended H to include $H(s)$.*

Proof. Suppose, for contradiction, that $G_1 \geq H_1$. Clearly, $G_1 \neq H_1$, because otherwise G and H would be the same bumping route and it would not be the case that the bump that extends G to include $G(r)$ came after the bump that extended H to include $H(r)$ for some r. Thus, $G_1 > H_1$, meaning that $H_1 < G_1$. By Theorem 3.21, this implies that the bump that extends H to include $H(r)$ came after the bump that extended G to include $G(r)$, a contradiction. Thus, $G_1 < H_1$. The rest of the corollary follows directly from Theorem 3.21. □

Corollary 3.24. *No point can be part of two different bumping routes that are created by the same application of multi-insertion. (Bumping routes of two different points are considered different, even if the points correspond to the same box.)*

Proof. Suppose, for contradiction, that this is false. Then there exists a point P that is part of two different bumping routes G and H. Without loss of generality, assume that $G_1 < H_1$ (we know that $G_1 \neq H_1$ because G and H are different bumping routes). Let $r(P)$ be the plane row that P is in. Then $G(r(P))$ and $H(r(P))$ are both defined. It follows by Theorem 3.21 that $H(r(P))$ is to the left of $G(r(P))$. However, by our assumption, $H(r(P)) = G(r(P)) = P$. This is a contradiction. Thus, no point can be part of two different bumping routes. □

Definition 3.25. Given any bumping route H, the *cylindric bumping route* $\pi(H)$ is the list of boxes $\pi(H_1)$, $\pi(H_2)$, ..., $\pi(H_{L(H)})$. Given any cylindric bumping route K, the set of bumping routes H such that $\pi(H) = K$ is denoted $\pi^{-1}(K)$.

Corollary 3.26. *No box is part of two different cylindric bumping routes. No box appears twice or more among the elements of any cylindric bumping route.*

Proof. Suppose for contradiction that a box B is part of two different cylindric bumping routes J and K. Let P be any element of $\pi^{-1}(B)$. Then there exists a bumping route $G \in \pi^{-1}(J)$ such that $P \in G$. Similarly, there exists a bumping route $H \in \pi^{-1}(K)$ such that $P \in H$. Clearly, G and H are distinct bumping routes, since $\pi(G) \neq \pi(H)$. This is a contradiction, by Corollary 3.24.

Suppose for contradiction that a box B appears twice or more among the elements of a cylindric bumping route K. Let x be the entry bumped out of B as a result of the bump that extends K to include B for the first time. Let y be the entry bumped out of B as a result of the bump that extends K to include B for the second time (we know that this is a

bump, not a landing, because B is already in the tableau); let z be the entry that bumps y out during this bump. We know that $z \geq x$, because bumped-out entries along a bumping route increase; we know that $y \leq x$ by Remark 3.4. It follows that $z \geq y$, a contradiction, as z bumps out y. Thus, no box appears twice or more among the elements of any cylindric bumping route, as desired. □

It follows from Corollary 3.26 that the entry that is in a particular box can only change once during an application of multi-insertion. This means that when an entry is inserted into a box, that entry remains in the box throughout the rest of the multi-insertion process. (In turn, this means that once an entry lands in a box, it cannot later be bumped from the box.) An entry cannot be bumped out more than once during an application of multi-insertion.

Corollary 3.27. *Let R be a tableau on which multi-insertion is performed. Then for all rows r of R, the list of entries inserted into r in the order in which they were inserted is weakly increasing.*

Proof. Suppose, for contradiction, that this is not the case. Then there exists a row r such that an entry is inserted into r, and then a smaller entry is inserted into r. Choose any $s \in \pi^{-1}(r)$. Let P and Q be points in s such that an entry a is inserted into P, and then $b < a$ is inserted into Q. Let G be the bumping route that contains P and H be the bumping route that contains Q (we know that G and H are unique from Corollary 3.24). From Corollary 3.23, we know that $H_1 < G_1$. From Theorem 3.21, we conclude that P is to the left of Q.

At the moment when Q becomes part of H, a has already been inserted into P, and from Corollary 3.26 we know that a stays in P for the rest of the insertion process. This means that b is inserted into Q while a larger entry a is in P, to its left. This is a contradiction. Therefore, for all rows r of R, the list of entries inserted into r in the order in which they were inserted is weakly increasing. □

Corollary 3.28. *Let G and H be two bumping routes. Then $G_1 < H_1$ if and only if $G_{L(G)} < H_{L(H)}$.*

Proof. We will first prove that if $G_1 < H_1$, then $G_{L(G)} < H_{L(H)}$. Suppose, for contradiction, that this is not the case — that there exist two bumping routes G and H such that $G_1 < H_1$, but $G_{L(G)} \geq H_{L(H)}$. It follows that the final point of G is in the same plane row as the final point of H, or it is in a lower plane row. If it is in the same plane row, it follows by Theorem 3.21 that $G_{L(G)}$ is to the right of $H_{L(H)}$, which means that $G_{L(G)} < H_{L(H)}$ — a contradiction.

If the final point of G is in a lower plane row than the final point of H, let r be the plane row such that $H_{L(H)} = H(r)$. Then $G(r)$ is defined, as otherwise it would not be the case that $G_1 < H_1$ (as then G would start lower than H). It follows by Theorem 3.21 that $H(r)$ is to the left of $G(r)$. However, since $H(r) = H_{L(H)}$, $\pi(H(r))$ is a box where an entry lands. This means that an entry lands in $\pi(G(r))$ as well. We know from Corollary 3.26 that an entry could not have landed in $\pi(G(r))$ and then have been bumped out. It follows that the insertion that extends G to include $G(r)$ is a landing. This is a contradiction, as then the final point of G cannot be in a lower plane row than the final point of H.

Now we will prove that if $G_{L(G)} < H_{L(H)}$, then $G_1 < H_1$. Suppose, for contradiction, that $G_1 \geq H_1$. We know that $G_1 \neq H_1$, because then G and H are the same bumping route, and

$G_{L(G)} = H_{L(H)}$. If $G_1 > H_1$, then, by the first part of this corollary (which we have proved), $H_{L(H)} < G_{L(G)}$. This is a contradiction. Therefore, if $G_{L(G)} < H_{L(H)}$, then $G_1 < H_1$. □

Definition 3.29. Let λ be the outer shape of a tableau R. Let S be a set of boxes such that R and S satisfy the input preconditions of Algorithm 3.12. Let $R' = \text{FullMulti}(R, S)$ and λ' be the outer shape of R'. Then the *new set* associated with R and S, denoted $N(R, S)$, is the set of boxes in λ'/λ.

Corollary 3.30. *For any tableau R and set of boxes S such that R and S are valid inputs into Algorithm 3.12, $N(R, S)$ forms a horizontal strip.*

Proof. Suppose, for contradiction, that $N(R, S)$ does not form a horizontal strip. Then there are two boxes $B, C \in N(R, S)$, with C immediately below B. Then there exist two points P and Q such that $P \in \pi^{-1}(B)$, $Q \in \pi^{-1}(C)$, and P is immediately above Q. Let G and H be the bumping routes that P and Q are a part of, respectively; thus, $P = G_{L(G)}$ and $Q = H_{L(H)}$. If $L(H) = 1$ (meaning that C was not in the tableau to begin with), P cannot be above Q, because then B would have had to be in λ to begin with, contradicting our supposition that $B \in N(R, S)$. If $L(H) > 1$, consider $Q' = H_{L(H)-1}$. We know that $\pi(P)$ is not in the outer shape of R to begin with. This means that $\pi(P)$ was added to R. By Corollary 3.26, once an entry has landed in $\pi(P)$, it is never bumped out. Thus, if $H_{L(H)-1}$ were P or any point directly to the right of P, Q would end at that landing point. This means that $H_{L(H)-1}$ is in the row of P, but to the left of P (and thus to the left of Q). This is a contradiction, as bumping routes trend weakly left. Thus, $N(R, S)$ necessarily forms a horizontal strip. □

4. Reverse Insertion and Reverse Multi-Insertion

4.1. Reverse Insertion Algorithms and Examples

Analogous to the process of row-deletion for non-cylindric tableaux, Algorithm 3.2 can be reversed in a process that we will call "reverse row-insertion." The algorithm for reverse row-insertion is shown below. We will prove later that reverse row-insertion is indeed a process that reverses internal row-insertion, and that the process always results in a valid semistandard tableau. The algorithm takes a tableau R and an outside corner B of R, defined as follows:

Definition 4.1. Given $R \in \text{SSCT}(\lambda/\mu)$, a box $B = \pi((i, j))$ is an *outside corner* of R if $B \in \lambda$, but $\pi((i + 1, j)) \notin \lambda$ and $\pi((i, j + 1)) \notin \lambda$. Visually, B is an outside cocorner of R if it is to the left of (inside) λ, but the boxes to the right of and below B are outside λ. Note that B is not necessarily a box of R, as it is possible that B is inside μ as well.

Algorithm 4.2 (Reverse Row-Insertion)**.**

Function ReverseInsert(tableau R, box B) ▷ B must be an outside corner of R.
1: $\mu :=$ inner shape of R.
2: $\lambda :=$ outer shape of R.

3: **if** $B \in R$ **then**:
4: $x :=$ entry of R that is in B.
5: **end if**.
6: Shrink λ to exclude B and remove B from R.
7: **if** $B \in \mu$ **then**: ▷ This happens only when B was not in R to begin with.
8: Shrink μ to exclude B.
9: **else**:
10: **while** $x \neq$ null **do**:
11: $r :=$ row of B.
12: **if** x is less than or equal to every entry in R in row $r-1$ **then**:
13: $B :=$ rightmost box of row $r-1$ that is in μ.
14: Put x in B and shrink μ to exclude B. ▷ We say that x *lands* in B.
15: $x :=$ null.
16: **else**:
17: $B :=$ box of the rightmost entry of R in row $r-1$ that is less than x.
18: $x' :=$ entry of R in B.
19: In R, replace x' in B with x.
20: $x := x'$.
21: **end if**.
22: **end while**.
23: **end if**.
24: **return** R.

For an example of reverse row-insertion, the reader may look at Example 3.3, starting with the last tableau drawn as R and the box with the 7 in the top row drawn as B. The tableaux obtained after each insertion are the same as the tableaux in Example 3.3, except in reverse order. The algorithm ends with the original tableau in the example.

Just as a particular box's entry can only decrease during insertion, the a particular box's entry can only increase during reverse insertion, since entries always bump out smaller entries.

Definition 4.3. Bumping routes for reverse insertion are defined exactly as bumping routes are defined for insertion. Consider a point P such that $\pi(P)$ was reverse-inserted into a tableau R. The *reverse bumping route* of P is a list of points, constructed as follows:

- Add P to the reverse bumping route when $\pi(P)$ is reverse-inserted into R.
- Say that a point $Q = (r, y)$ is added to the reverse bumping route when an entry x is bumped out of $\pi(Q)$. When x is reverse-inserted into a box $B \in \pi(r-1)$, add the element of $\pi^{-1}(B)$ that is in plane row $r-1$ to the bumping route.

Since x decreases after every iteration of the loop beginning on line 10 and x cannot be smaller than the smallest entry in the original tableau, the loop must terminate and the algorithm must end. Thus, reverse insertion, just like insertion, always ends after a finite number of steps.

Definition 4.4. An insertion queue is *reverse-regular* if, for any two elements of the queue (x_1, r) and (x_2, r), where $x_1 < x_2$, (x_1, r) comes *after* (x_2, r) in the queue.

Just as with insertion, one can reverse-insert multiple entries at the same time. The algorithm below is analogous to Algorithm 3.9; it takes a tableau and a reverse-regular insertion queue as input and outputs a pair consisting of a map from a subset of $\mathcal{C}_{k,n}$ to the alphabet of the tableau taken as input (though the map itself is not necessarily a valid semistandard tableau) and an insertion queue.

Algorithm 4.5 (Reverse One-Step Multi-Insertion)**.**

Function ReverseOneStepMulti(tableau R, insertion queue $\mathfrak{q}$) ▷ $\mathfrak{q}$ must be reverse-regular.
1: $\mu :=$ inner shape of R.
2: $\mathfrak{q}' :=$ empty insertion queue.
3: **while** $\mathfrak{q}$ is not empty **do**:
4: Remove the first element from $\mathfrak{q}$. Let it be (x, r).
5: **if** x is less than or equal to every entry in R in row r **then**:
6: $B :=$ rightmost box of row r that is in μ.
7: Put x in B and shrink μ to exclude B. ▷ μ is not necessarily a valid partition anymore.
8: **else**:
9: $B :=$ box of the rightmost entry of R in row r that is less than x.
10: $x' :=$ entry of R in B.
11: In R, replace x' in B with x.
12: Add $(x', r-1)$ to $\mathfrak{q}'$.
13: **end if**.
14: **end while**.
15: **return** $(R, \mathfrak{q}')$. ▷ R is not necessarily a tableau.

Remark 4.6. *The insertion queue returned by Algorithm 4.5 is reverse-regular.*

Proof. Let $(R', \mathfrak{q}') =$ ReverseOneStepMulti$(R, \mathfrak{q})$ for a tableau R and a reverse-regular insertion queue $\mathfrak{q}$. Suppose, for contradiction, that $\mathfrak{q}'$ is not reverse-regular. Then there exist two elements of $\mathfrak{q}'$, (y_1, r) and (y_2, r), such that $y_1 > y_2$, but (y_2, r) comes before (y_1, r) in $\mathfrak{q}'$. Let x_1 and x_2 be the entries that bumped out y_1 and y_2, respectively. Then $(x_2, r+1)$ and $(x_1, r+1)$ were in $\mathfrak{q}$, with $(x_2, r+1)$ coming first. Since $\mathfrak{q}$ is reverse-regular, it follows that $x_2 > x_1$. We also know that $x_1 > y_1$, since x_1 bumps out y_1. It follows that $x_2 > x_1 > y_1 > y_2$. Since $y_1 < y_2$, y_1 is to the right of y_2 in R, and $y_1 < x_2$. It follows that y_2 is not the rightmost entry of R in its row that is greater than x_2 at the time that x_2 is to be inserted, which is a contradiction, because then x_2 does not bump out y_2. Thus, $\mathfrak{q}'$ is a reverse-regular insertion queue. □

We can fully reverse-insert multiple entries from a tableau simultaneously using Algorithm 4.5 as a subroutine, in a way analogous to Algorithm 3.12. This algorithm takes a tableau and a set of boxes that forms a horizontal strip as input (with the precondition that, when the outer shape of the tableau is shrunk to exclude these boxes, it remains a valid partition) and outputs a tableau (we will prove later that the output is indeed a valid semistandard tableau).

Algorithm 4.7 (Reverse Full Multi-Insertion)**.**

Function ReverseFullMulti(tableau R, set S of boxes) ▷ The outer shape of R minus the boxes in S must be a valid partition; every box in S is in the outer shape of R. Also, S must form a horizontal strip.

$\mu :=$ inner shape of R.
$\lambda :=$ outer shape of R.
$\mathfrak{q}_0 :=$ empty insertion queue.
Choose any integer r. ▷ We prove later that the choice of r is immaterial.
$h := r$.
while $h \neq r - k$ **do**: ▷ k is the vertical period of R.
    $L :=$ list of boxes in row r in S, from right to left.
    while L is not empty **do**:
        $B :=$ first element of L.
        Remove B from L.
        if $B \in R$ **then**:
            $x :=$ entry of R in B.
            Put $(x, h-1)$ into $\mathfrak{q}_0$.
            Shrink λ to exclude B and remove B from R.
        else:
            Shrink μ and λ to exclude B and remove B from R.
        end if.
    end while.
    $h := h - 1$.
end while.
$R_0 := R$.
$i := 0$.
while $\mathfrak{q}_i$ is not empty **do**:
    $(R_{i+1}, \mathfrak{q}_{i+1}) :=$ ReverseOneStepMulti $(R_i, \mathfrak{q}_i)$. ▷ This line inserts all elements of $\mathfrak{q}_i$ into R_i and calls the result R_{i+1}. $\mathfrak{q}_{i+1}$ is the insertion queue of bumped-out cells.
    $i := i + 1$.
end while.
$R := R_i$.
return R.

Example 4.8. Let R, drawn below, be the input tableau into Algorithm 4.7, and let the red boxes in the diagram below constitute S.

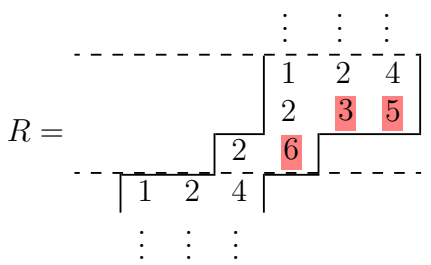

Suppose that we pick row 1 to be our starting row r. We shrink λ to exclude the red boxes, deleting those boxes from R and putting corresponding entries into $\mathfrak{q}_0$. The 5 is deleted from row 1, and it is to be reverse-inserted into the previous row, so we add $(5,0)$ to $\mathfrak{q}_0$. We then add $(3,0)$ to $\mathfrak{q}_0$ ($(5,0)$ is added first because entries are processed from right to left, not from left to right as done in Algorithm 3.12). Finally, $(6,1)$ is added to $\mathfrak{q}_0$. We let R_0 be our current tableau.

$$\mathfrak{q}_0 = (5,0),(3,0),(6,1) \qquad R_0 =$$

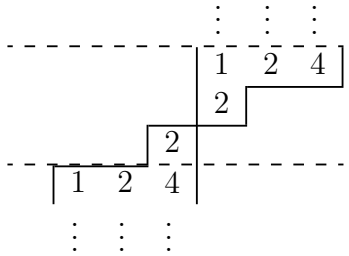

We now enter the One-Step Reverse Multi-Insertion subroutine. We remove $(5,0)$ from $\mathfrak{q}$ and insert 5 into row 0, bumping out the 4 (the rightmost entry strictly less than 5) and adding $(4,2)$ to $\mathfrak{q}'$. We remove $(3,0)$ from $\mathfrak{q}$ and insert 3 into row 0, bumping out the 2 and adding $(2,2)$ to $\mathfrak{q}'$. We remove $(6,1)$ from $\mathfrak{q}$ and insert 6 into row 1, bumping out the 2 and adding $(2,0)$ to $\mathfrak{q}'$. Finally, with $\mathfrak{q}$ empty, we exit the subroutine. We let $\mathfrak{q}_1$ and R_1 be the returned insertion queue and the returned tableau, respectively; they are shown below:

$$\mathfrak{q}_1 = (4,2),(2,2),(2,0) \qquad R_1 =$$

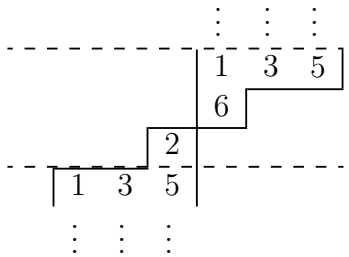

The process continues as follows:

$$\mathfrak{q}_2 = (2,1),(1,2) \qquad R_2 =$$

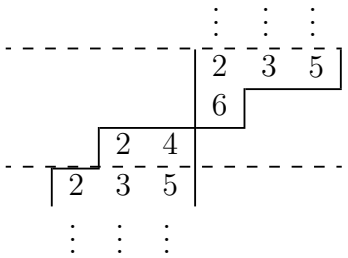

$$\mathfrak{q}_3 \text{ is empty} \qquad R_3 =$$

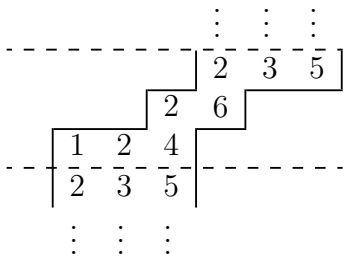

Now that $\mathfrak{q}_3$ is empty, we exit out of the loop, let R equal R_3, and return R.

Since entries only bump out smaller entries and no bumped-out entry can be smaller than the smallest entry originally in the tableau, any chain of bumps eventually ends. It follows that Algorithm 4.7 always terminates.

In the same way that the concept of a bumping route applies to multi-insertion as well as single-box insertion, the concept of a reverse bumping route applies to reverse multi-insertion as well as reverse single-box insertion.

4.2. Relating Reverse Insertion to Forward Insertion

Theorem 4.9. *Let R be a tableau and S be a set of boxes such that R and S satisfy the input preconditions of Algorithm 3.12. Let $R' = \text{FullMulti}(R, S)$ and $S' = N(R, S)$ (see Definition 3.29). Let B be a box in S. Choose $P \in \pi^{-1}(B)$ and let H be the bumping route of P. Let $P' = H_{L(H)}$ and let H' be the reverse bumping route of P' when Algorithm 4.7 is called on R' and S'. Then, for all $1 \leq i \leq L(H)$, $H'_i = H_{L(H)+1-i}$. That is, H' consists of the same points as H, but in reverse order.*

Proof. We first note that the preconditions for Algorithm 4.7 are satisfied: S' forms a horizontal strip (Corollary 3.30) that can be deleted from the outer shape of R'.

Take any R and S as described in the theorem statement. We know by definition that, for all $P \in \bigcup_{B \in S} \pi^{-1}(B)$ (with P', H, and H' defined correspondingly as in the theorem statement), $H'_1 = P' = H_{L(H)}$. Suppose that, up to but not including the reverse-insertion of an entry x, all reverse bumping routes created by the reverse multi-insertion process have retraced the bumping routes created by the multi-insertion process. Let A be the box containing x prior to this reverse insertion. Select any $Q \in \pi^{-1}(A)$. Let H be the bumping route containing Q that was made by the multi-insertion algorithm and H' be the reverse bumping route containing Q that was made by the reverse multi-insertion algorithm. Let j be such that $H'_j = Q$. (Thus, $H'_j = H_{L(H)+1-j}$.) We will show that $H'_{j+1} = H_{L(H)-j}$.

Let r be the plane row immediately above the plane row containing Q. Let C be the box into which x is to be reverse-inserted if $H'_{j+1} = H_{L(H)-j}$. Suppose, for contradiction, that x is not reverse-inserted into C.

We know that x was bumped out of C in the forward multi-insertion process. Let y be the entry that replaced x in C. By Corollary 3.26, we know that x and y were the only two entries ever to be present in C during the forward multi-insertion process. Since all bumping routes thus far have been retraced, it follows that no entry has thus far been reverse-inserted into C. Thus, y is in C when x is to be reverse-inserted into $\pi(r)$. Since $y < x$ and x is not reverse-inserted into C, it follows that y is not the rightmost entry in $\pi(r)$ (or r) that is less than x. Let z be the rightmost entry in r that is less than x, and let z be in box D.

Let V be the box in plane row r onto which C projects. Let W be the box in plane row r onto which D projects. We know that W contains an entry less than x; however, this was not the case when y was inserted into C during the forward multi-insertion process (because then z, which is smaller than x, would have been to the right of x). Since the entry in D changes, W must be part of a bumping route. Call this bumping route G, and let G' be the corresponding reverse bumping route to G created during the reverse multi-insertion process. Since z is still in W by the time that x is to be inserted into plane row r, the first point of

G' must either be below the first point of H' (in which case G' reaches plane row r on a later call of the subroutine), or be to the left of the first point of H' (in which case it comes later in the queue and continues to be later in the queue, because of the way that the elements of the queue are processed). This means that $H_{L(H)} = H'_1 < G'_1 = G_{L(G)}$.

Since V is to the left of W, H_r is to the left of G_r. By Corollary 3.22, it follows that $G_1 < H_1$. From this, by Corollary 3.28, it follows that $G_{L(G)} < H_{L(H)}$. This is a contradiction. Therefore, x is reverse-inserted into C.

Having completed our induction, we have proven that all reverse bumping routes retrace their corresponding bumping routes during reverse multi-insertion. □

Remark 4.10. *Since all bumping routes are retraced, multi-insertion followed by reverse multi-insertion results in the original tableau. That is, for any R and S satisfying the pre-conditions of multi-insertion,*

$$\text{ReverseFullMulti}(\text{FullMulti}(R, S), N(R, S)) = R.$$

By letting S consist of one box, we see that Algorithm 4.2 retraces the bumping route of Algorithm 3.2 and returns the original tableau. This means that everything henceforth that we prove about reverse multi-insertion also applies to reverse insertion of a single box.

Definition 4.11. The *negation function* Neg takes a tableau R and, for every entry x in R, changes x to $\overline{x}$. For any two entries x and y such that $x < y$, we have $\overline{y} < \overline{x}$. If the entries of our tableau are integers, we can think of the Neg operator as taking the additive inverse of every entry in the tableau and returning the result.

Definition 4.12. Define the *flip function* on a cylindric partition λ as follows: $\text{Flip}(\lambda)$ is the partition such that a point (x, y) is in $\text{Flip}(\lambda)$ if and only if it is not in λ. Equivalently, $\text{Flip}(\lambda)$ is the cylindric partition with sequence $(\dots, -1 - \lambda_1, -1 - \lambda_0, -1 - \lambda_{-1}, \dots)$.

Consider the map $M : \mathbb{Z}^2 \to \mathbb{Z}^2$ that takes (x, y) to $(-x, -y)$. Since M takes $(-k, n-k)\mathbb{Z}$ to itself, it follows that M induces an endomorphism E of $\mathcal{C}_{k,n}$. If we think of a cylindric tableau as a partial map from the cylinder to a totally ordered set, then, for a cylindric tableau R, define $\text{Flip}(R)$ to be the map created by applying E on the cylinder, followed by R, followed by Neg. If R has inner and outer shapes μ and λ, respectively, then $\text{Flip}(R)$ has inner and outer shapes $\text{Flip}(\lambda)$ and $\text{Flip}(\mu)$, respectively.

We also define the Flip function for boxes and sets of boxes. If B is a box, then $\text{Flip}(S)$ is the endomorphism E described above applied to B. If S is a set of boxes, $\text{Flip}(S)$ is E applied to S.

Simply put, let R be a tableau and B be a box of R that contains t. Choose any $P \in \pi^{-1}(B)$ and let P have coordinates (a, b). Let $C = \pi((-a, -b))$. Then $C = \text{Flip}(B)$ and, in the tableau $\text{Flip}(R)$, C contains $\overline{t}$.

Visually, Flip takes a tableau, rotates it by 180°, and negates all of the tableau's entries. Similarly, Flip takes a set of boxes and rotates it by 180°.

Remark 4.13. *Let R be a tableau. Then* $\text{Flip}(R)$ *is also a tableau.*

Proof. Clearly, $\text{Flip}(R)$ is bounded by two valid partitions — a partition remains valid after a 180° rotation; equivalently, Flip negates all elements of a partition's sequence, but also reverses their order, keeping the sequence weakly decreasing.

Let B and C be two boxes in $\text{Flip}(R)$ such that C is directly to the right of B. Then $\text{Flip}(B)$ is directly to the right of $\text{Flip}(C)$, so, in R, $\text{Flip}(B)$ contains an entry that is weakly greater than the entry in $\text{Flip}(C)$. This means that, in $\text{Flip}(R)$, B contains an entry that is weakly less than the entry in C, as desired.

Similarly, let D and E be two boxes in $\text{Flip}(R)$ such that D is directly below E. Then $\text{Flip}(D)$ is directly above $\text{Flip}(E)$, so, in R, $\text{Flip}(E)$ contains an entry that is strictly greater than the entry in $\text{Flip}(D)$. This means that, in $\text{Flip}(R)$, E contains an entry that is strictly less than the entry in D, as desired. Therefore, $\text{Flip}(R)$ is a valid semistandard tableau. □

Remark 4.14. Flip *is an involution — that is,* $\text{Flip} \circ \text{Flip} = \text{I}$, *where* I *is the identity function.*

Remark 4.15. *Let S be a set of boxes that forms a horizontal strip. Then* $\text{Flip}(S)$ *forms a horizontal strip.*

Theorem 4.16. *Let R be a tableau and S be a set of boxes such that R and S satisfy the input preconditions of Algorithm 4.7. Call Algorithm 4.7 on R and S; simultaneously, call Algorithm 3.12 on $R' = \text{Flip}(R)$ and $S' = \text{Flip}(S)$. Then $R'_0 = \text{Flip}(R_0)$. Furthermore, for every value of t, the tableau in Algorithm 3.12 after the t'th insertion is the flip of the tableau in Algorithm 3.12 after the t'th insertion; the same number of insertions occurs in the two algorithms.*

Proof. We first note that the preconditions for Algorithm 3.12 are satisfied: $\text{Flip}(R)$ is a valid semistandard tableau (Remark 4.13) and $\text{Flip}(S)$ forms a horizontal strip (Remark 4.15) that can be added to the inner shape of $\text{Flip}(R)$ and keep the partition valid. We will show our theorem to be true by simultaneously performing steps of the reverse multi-insertion algorithm on R and S and the forward multi-insertion algorithm on R' and S'.

Clearly, $R'_0 = \text{Flip}(R_0)$, as R_0 is R without the boxes in S, and R'_0 is $\text{Flip}(R)$ without the boxes in $\text{Flip}(S)$.

Let r_1 be the value of r chosen in line 4 of Algorithm 4.7. Then choose $-r_1$ for the value of r in line 4 of Algorithm 3.12.

Definition 4.17. The *corresponding element* of an element (x, r) of an insertion queue is $(\overline{x}, -r)$. Two insertion queues $\mathfrak{q}_1$ and $\mathfrak{q}_2$ are *corresponding queues* is they contain the same number of elements and, for all i, the i'th elements of $\mathfrak{q}_1$ and $\mathfrak{q}_2$ are corresponding elements.

Let $\mathfrak{q}_0$ be the original queue of elements taken out of R during the reverse multi-insertion process. Let $\mathfrak{q}'_0$ be the original queue of elements taken out of R' during the forward multi-insertion process. We first note that $\mathfrak{q}_0$ and $\mathfrak{q}'_0$ are corresponding queues. Since $\mathfrak{q}_0$ contains pairs containing all entries in the boxes of S in R and $\mathfrak{q}'_0$ contains pairs containing all entries in the boxes of S' in R', and these entries are in rows whose numbers are negatives of one another, we know that the elements of $\mathfrak{q}_0$ and $\mathfrak{q}'_0$ can be paired such that each element is paired with its corresponding element. In order to show that these entries are in the same order, consider two boxes $B, C \in S$ and their corresponding boxes $B' = \text{Flip}(B)$ and $C' = \text{Flip}(C)$. If B and C are in the same row, suppose, without loss of generality, that B is to the left of C. Then C is deleted from R (and its entry is put into $\mathfrak{q}_0$) before B is deleted from R. B' is to the right of C', so the entry in C' is put in a pair that is put into

$\mathfrak{q}_0'$ before the entry in B' is put into a pair that is put into $\mathfrak{q}_0'$. Thus, C comes before B and C' comes before B' as desired. If B and C are in different rows, suppose, without loss of generality, that B is deleted from R before C is. Since rows are processed from bottom to top beginning with row r_1 in the reverse multi-insertion process and they are processed from top to bottom beginning with row $-r_1$ in the forward multi-insertion process (and the rows of B' and C' are opposite to those of B and C, respectively), we conclude that B' is deleted from R' before C' is. Thus, for any two entries in $\mathfrak{q}_0$, their corresponding entries in $\mathfrak{q}_0'$ are in the same order relative to each other as are the two entries in $\mathfrak{q}_0$. We conclude that $\mathfrak{q}_0$ and $\mathfrak{q}_0'$ are corresponding queues.

We have shown that, when the respective subroutines are called on R_0 and R_0', the tableaux are flips of each other and the queues referred to as $\mathfrak{q}$ by the subroutines are corresponding. Since the queues referred to as $\mathfrak{q}'$ by the subroutines are initially empty, they are also corresponding. Suppose that we move through both algorithms step by step, examining the tableaux and queues after every insertion. Suppose that, up to but not including the insertion of x into row s in the reverse insertion process, the queues referred to as $\mathfrak{q}$ in the subroutines are corresponding, the queues referred to as $\mathfrak{q}'$ in the subroutines are corresponding, and the two tableaux are flips of each other. Let V be the tableau in the reverse insertion process right before the insertion of x into row s and V' be the corresponding tableau in the forward insertion process (thus, $V' = \text{Flip}(V)$). We will show that the two pairs of queues continue to correspond and that $V' = \text{Flip}(V)$ after the insertion as well.

Suppose that x takes the place of y in row s in V. Then y is the rightmost entry smaller than x in s. By our inductive hypothesis, in V', $\overline{y}$ is in row $-s$, and, by definition, $\overline{y} > \overline{x}$. Suppose that there is an entry left of $\overline{y}$ in V' that is greater than $\overline{x}$; call this entry $\overline{z}$. Then in V, $z < x$ and z is to the right of y; this is a contradiction. Thus, $\overline{x}$ takes the place of $\overline{y}$ in row $-s$ in V'. The first entries in the queues called $\mathfrak{q}$ by the subroutines are deleted, so these queues continue to correspond. Corresponding entries ($(y, s-1)$ and $(\overline{y}, -s+1)$) are added to the queues referred to as $\mathfrak{q}'$, so these queues also continue to correspond. Thus, our inductive step holds, as desired.

We have completed our induction. Since both algorithms end when their respective queues (which are corresponding and thus have the same length) are empty, the two algorithms terminate after the same total number of insertions, and, after each successive insertion, the two tableaux remain flips of one another. □

Definition 4.18. Let μ be the outer shape of a tableau R. Let S be a set of boxes such that R and S satisfy the preconditions of Algorithm 4.7. Let $R' = \text{ReverseFullMulti}(R, S)$ and μ' be the outer shape of R'. Then the *reverse new set* associated with R and S, denoted $N^{-1}(R, S)$, is the set of boxes in μ/μ'.

We will now create alternative inputs and outputs for the functions FullMulti and ReverseFullMulti. Each algorithm can now take a pair (R, S) (instead of two arguments R and S) and returns a pair (R', S'), the first element being the tableau generated by the algorithms and the second being:

- $S' = N(R, S)$ for FullMulti; and
- $S' = N^{-1}(R, S)$ for ReverseFullMulti.

We will also define $\text{Flip}((R, S))$ to be $(\text{Flip}(R), \text{Flip}(S))$.

Corollary 4.19.

$$\text{ReverseFullMulti}((R, S)) = \text{Flip}(\text{FullMulti}(\text{Flip}((R, S)))).$$

Proof. Given R and S, since, after all insertions have been accounted for, the tableaux returned by Algorithm 4.7 and Algorithm 3.12 are flips of one another, the tableaux returned by two algorithms must be flips of one another. Thus, flipping the tableau returned by Algorithm 3.12 results in the tableau returned by Algorithm 4.7. Since $\text{FullMulti}(\text{Flip}((R, S)))$ is the flip of $\text{ReverseFullMulti}((R, S))$, it follows that $N^{-1}(R, S)$ is the flip of $N(\text{Flip}((R, S)))$, as desired. □

We can use Theorem 4.16 to prove about reverse multi-insertion much of what we proved about multi-insertion.

Remark 4.20. *Given R and S satisfying the preconditions of Algorithm 4.7, R_0 is a valid semistandard tableau. Furthermore, after every insertion that occurs while Algorithm 4.7 is being performed (including during the subroutine), R is a valid semistandard tableau.*

Proof. We know from Theorem 4.16 that any value of R that appears during the reverse multi-insertion process, starting from R_0, is the flip of a value of R that appears during the forward multi-insertion process when it is applied to $\text{Flip}(R)$ and $\text{Flip}(S)$, starting from $\text{Flip}(R)_0$. We know that all values of R starting from R_0 during forward multi-insertion are indeed valid semistandard tableaux from Proposition 3.16. Since the Flip operator preserves valid semistandard tableaux (Remark 4.13), R is a valid semistandard tableau after every insertion, as desired. □

Remark 4.21. *The choice of r in line 4 of Algorithm 4.7 is irrelevant.*

Proof. We have shown that $\text{ReverseFullMulti}(R, S) = \text{Flip}(\text{FullMulti}(\text{Flip}(R), \text{Flip}(S)))$, where the r chosen in the forward insertion process is the opposite of the r chosen in the reverse insertion process. However, the r chosen in the forward insertion process is immaterial (Remark 3.18). Thus, the choice of r in the reverse insertion process cannot matter either. □

Remark 4.22. *If $(R', S') = \text{ReverseFullMulti}((R, S))$, then S' forms a horizontal strip.*

Proof. By Corollary 4.19, we know that

$$\text{ReverseFullMulti}((R, S)) = \text{Flip}(\text{FullMulti}(\text{Flip}((R, S)))).$$

Since the Flip function preserves horizontal strips (Remark 4.15) and FullMulti returns a set of boxes that forms a horizontal strip (Corollary 3.30), we conclude that $\text{Flip}(\text{FullMulti}(\text{Flip}((R, S))))$ returns a pair whose second element does indeed form a horizontal strip. □

Remark 4.23. *Let R be a tableau and S a set of boxes on which reverse multi-insertion is performed. Then for all rows r of R, the list of entries bumped out of r in the order in which they were bumped out is weakly decreasing.*

Proof. Let $R' = \text{Flip}(R)$ and $S' = \text{Flip}(S)$. By Theorem 4.16, the reverse insertion tableau is the flip of the forward insertion tableau after every time that an insertion is completed in both algorithms. Also, during the formation of R_0 and R'_0, entries are bumped out in the same order. This means that if $L = a_1, a_2, \ldots, a_m$ is the list of entries bumped out of a row r during the reverse multi-insertion process, the list of entries bumped out of $-r$ during the forward multi-insertion process at the same point during the process is $\overline{a_1}, \overline{a_2}, \ldots, \overline{a_m}$.

By Lemma 3.19, $\overline{a_1}, \overline{a_2}, \ldots, \overline{a_m}$ is a weakly increasing list. Therefore, $L = a_1, a_2, \ldots, a_m$ is a weakly decreasing list, as desired. □

Analogously, we can use Corollary 3.27 to prove that, given a row r of a tableau R on which reverse multi-insertion is performed, the list of entries inserted into r in the order in which they were inserted is weakly decreasing.

Proposition 4.24. ReverseFullMulti *and* FullMulti *are inverse operations.*

Proof. In order to make notation less cumbersome, we will shorten the names of the operations as follows:

- ReverseFullMulti to RFM;
- FullMulti to FM; and
- Flip to F.

We have shown that $\text{RFM}(\text{FM}((R, S))) = (R, S)$ (Remark 4.10) — that is, $\text{RFM} \circ \text{FM} = \text{I}$. (The fact that the second element of the output pair is S follows from the fact that the boxes added back to R after being removed are the ones that are originally removed.) It remains to show only that $\text{FM} \circ \text{RFM} = \text{I}$.

Consider the following facts:

- From Remark 4.10:
$$\text{RFM} \circ \text{FM} = \text{I} \tag{1}$$
- From Remark 4.14:
$$\text{F} \circ \text{F} = \text{I} \tag{2}$$

We proceed as follows:

Adding F to the left of both sides of (1) and applying (2), we have

$$\text{F} \circ \text{RFM} = \text{F} \circ \text{F} \circ \text{FM} \circ \text{F} = \text{I} \circ \text{FM} \circ \text{F} = \text{FM} \circ \text{F}\,.$$

Adding FM to the right of both sides, we have

$$\text{FM} \circ \text{F} \circ \text{FM} = \text{F} \circ \text{RFM} \circ \text{FM} = \text{F} \circ \text{I} = \text{F}\,. \tag{3}$$

Now adding F to the right of both sides of (1) and applying (2), we have

$$\text{RFM} \circ \text{F} = \text{F} \circ \text{FM} \circ \text{F} \circ \text{F} = \text{F} \circ \text{FM} \circ \text{I} = \text{F} \circ \text{FM}\,.$$

Adding FM to the left of both sides (here we need Remark 4.22 for well-definedness), we have

$$\mathrm{FM} \circ \mathrm{RFM} \circ \mathrm{F} = \mathrm{FM} \circ \mathrm{F} \circ \mathrm{FM}\,.$$

Applying (3), we have

$$\mathrm{FM} \circ \mathrm{RFM} \circ \mathrm{F} = \mathrm{FM} \circ \mathrm{F} \circ \mathrm{FM} = \mathrm{F}\,.$$

Adding F to the right of both sides, we have

$$\mathrm{FM} \circ \mathrm{RFM} \circ \mathrm{F} \circ \mathrm{F} = \mathrm{F} \circ \mathrm{F}\,.$$

Canceling out the F's, we have

$$\mathrm{FM} \circ \mathrm{RFM} = \mathrm{I},$$

as desired. □

Having proven this, we have created a bijection between pairs consisting of a cylindric tableau and a set of boxes that forms a horizontal strip that can be added to the inner shape of the tableau while keeping the partition valid on the one hand, and pairs consisting of a cylindric tableau and a set of boxes that forms a horizontal strip that can be removed from the outer shape of the tableau while keeping the partition valid on the other. This is the basis of the cylindric Robinson-Schensted-Knuth (RSK) correspondence, which will be detailed in the following section.

4.3. More Results about Reverse Insertion

We can use the bijection proven above in order to prove more results about reverse multi-insertion, particularly about bumping routes. This is because, while we proved earlier that forward bumping routes are retraced with reverse multi-insertion (Theorem 4.9), before this Proposition 4.24 we did not prove that FullMulti(ReverseFullMulti$((R, S))) = (R, S)$. This result allows us to use more extensively what we already know about forward multi-insertion in order to prove results about reverse multi-insertion.

Remark 4.25. *Reverse bumping routes trend weakly right — that is, given a reverse bumping route G generated by performing reverse multi-insertion on (R, S), $G(m)$ is to the right of G_{m-1} for all $1 < m \leq L(G)$.*

Proof. Let $(R', S') =$ ReverseFullMulti$((R, S))$. We know that FullMulti$((R', S')) = (R, S)$. Let G' be the bumping route that consists of the same points as G (we know such a G' exists from Theorem 4.9). G' trends weakly left, and G has the same points as G', but in reverse order. Therefore, G trends weakly right, as desired. □

Remark 4.26. *(Counterpart of Theorem 3.21.) Let R be a tableau and S be a set of boxes such that R and S satisfy the input preconditions for Algorithm 4.7. Let G and H be two reverse bumping routes created when Algorithm 4.7 is performed with R and S as input, such that $G_1 < H_1$. Then for all plane rows r such that $G(r)$ and $H(r)$ are defined, $H(r)$ is strictly to the left of $G(r)$ and the bump that extended G to include $G(r)$ came* before *the bump that extended H to include $H(r)$.*

Proof. Let G' and H' be the forward bumping routes that correspond to G and H, respectively (obtained by performing Algorithm 3.12 on the result of ReverseFullMulti((R, S))). We know that $G'_{L(G')} = G_1 < H_1 = H'_{L(H')}$, so, by Corollary 3.28, $G'_1 < H'_1$. By Theorem 3.21, this means that, for all plane rows r such that $G'(r)$ and $H'(r)$ are defined, $H'(r)$ is strictly to the left of $G'(r)$. Since $H'(r) = H(r)$ and $G'(r) = G(r)$ for all r, it follows that $H(r)$ is strictly to the left of $G(r)$.

Let s be any number such that $G(s)$ and $H(s)$ are defined. If G_1 and H_1 are in different plane rows, then G_1 is above H_1, which means that the bump that extended G to include $G(s)$ came in an earlier call of the subroutine than the bump that extended H to include $H(s)$. If G_1 and H_1 are in the same plane row, let t be the plane row of G_1 and H_1 and let $u = \pi(t)$. Let a and b be the entries in $\pi(G_1)$ and $\pi(H_1)$, respectively. Since $G_1 < H_1$, $\pi(G_1)$ is to the right of $\pi(H_1)$, which means that $(b, u-1)$ is inserted into $\mathfrak{q}_0$ before $(a, u-1)$ is. Because of how entries are taken out of and put into new queues in the subroutine, the pair corresponding to G comes before the pair corresponding to H in $\mathfrak{q}_m$ for all m. Thus, in this case as well, the bump that extended G to include $G(s)$ came before the bump that extended H to include $H(s)$. Therefore, for all plane rows r such that $G(r)$ and $H(r)$ are defined, the bump that extended G to include $G(r)$ came before the bump that extended H to include $H(r)$. □

Remark 4.27. *(Counterpart of Corollary 3.22.) Let G and H be two reverse bumping routes such that, for some plane row r, $H(r)$ is strictly to the left of $G(r)$. Then $G_1 < H_1$, and for all plane rows s such that $G(s)$ and $H(s)$ are defined, $H(s)$ is strictly to the left of $G(s)$.*

Proof. Let G' and H' be the forward bumping routes corresponding to G and H. Since $H(r)$ is strictly to the left of $G(r)$, $H'(r)$ is strictly to the left of $G(r)$. By Corollary 3.22, this implies that, for all plane rows s such that $G'(s)$ and $H'(s)$ are defined, $H'(s)$ is strictly to the left of $G'(s)$, which implies that, for all plane rows s such that $G(s)$ and $H(s)$ are defined, $H(s)$ is strictly to the left of $G(s)$. It also implies that $G'_1 < H'_1$, which, by Corollary 3.28, implies that $G'_{L(G')} < H'_{L(H')}$, which implies that $G_1 < H_1$, as desired. □

Remark 4.28. *(Counterpart of Corollary 3.23; the proof is entirely analogous, so it is not given here.) Let G and H be two reverse bumping routes such that, for some plane row r, the bump that extends G to include $G(r)$ came after the bump that extended H to include $H(r)$. Then $G_1 < H_1$, and for all plane rows s such that $G(s)$ and $H(s)$ are defined, the bump that extends G to include $G(s)$ came before the bump that extended H to include $H(s)$.*

Remark 4.29. *(Counterpart of Corollary 3.24.) No point can be part of two different reverse bumping routes that are created by the same application of multi-insertion.*

Proof. Suppose that a point P is part of two reverse bumping routes G and H. Let G' and H' be the corresponding forward bumping routes to G and H. Then P is on both G and H. By Corollary 3.24, this is a contradiction. □

Definition 4.30. Given any reverse bumping route H, the *cylindric reverse bumping route* $\pi(H)$ is the list of boxes $\pi(H_1)$, $\pi(H_2)$, ..., $\pi(H_{L(H)})$. Given any cylindric reverse bumping route K, the set of bumping routes H such that $\pi(H) = k$ is denoted $\pi^{-1}(K)$.

Remark 4.31. *(Counterpart of Corollary 3.26.) No box is part of two different cylindric reverse bumping routes. No box appears twice of more among the elements of a cylindric bumping route.*

Proof. Suppose there is a box B that is part of two different cylindric reverse bumping routes J and K. Let P be any element of $\pi^{-1}(B)$. Then there exists a reverse bumping route $G \in \pi^{-1}(J)$ and a reverse bumping route $H \in \pi^{-1}(K)$ such that $P \in G$ and $P \in H$. By Remark 4.29, this is a contradiction.

Suppose that a box B appears twice of more among the elements of a cylindric reverse bumping route K. Choose any $H \in \pi^{-1}(K)$. Let H' be the corresponding forward bumping route to H. Then $\pi(H')$ contains B twice or more. By Corollary 3.26, this is a contradiction. □

Remark 4.32. *(Counterpart of Corollary 3.28.) Let G and H be two reverse bumping routes. Then $G_1 < H_1$ if and only if $G_{L(G)} < H_{L(H)}$.*

Proof. Let G' and H' be the corresponding forward bumping routes to G and H, respectively. We know that $G_1 = G'_{L(G')}$, $H_1 = H'_{L(H')}$, $G_{L(G)} = G'_1$, and $H_{L(H)} = H'_1$. From Corollary 3.28, we know that $G'_1 < H'_1$ if and only if $G'_{L(G')} < H'_{L(H')}$. Therefore, $G_1 < H_1$ if and only if $G_{L(G)} < H_{L(H)}$. □

5. The Cylindric RSK Correspondence

5.1. The Correspondence

We now present an analog of the Robinson-Schensted-Knuth correspondence for cylindric tableaux. The idea for this correspondence is based on that of Sagan and Stanley [SagStan, §6, Theorem 6.1].

Definition 5.1. Let S be a set and T_s be a set for any $s \in S$. We define the *disjoint union* operator $\bigsqcup$ as follows:

$$\bigsqcup_{s \in S} T_s = \{(t, s) | s \in S, t \in T_s\}$$

Theorem 5.2. *Fix two partitions α and β. Then there exists a bijective mapping*

$$\text{CRSK} : \bigsqcup_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha;\ \mu \subseteq \beta}} \text{SSCT}(\alpha/\mu) \times \text{SSCT}(\beta/\mu) \to \bigsqcup_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda;\ \beta \subseteq \lambda}} \text{SSCT}(\lambda/\beta) \times \text{SSCT}(\lambda/\alpha)$$

such that for any $\mu \in$ Cylpar satisfying $\mu \subseteq \alpha$ and $\mu \subseteq \beta$, $T \in \text{SSCT}(\alpha/\mu)$, and $U \in \text{SSCT}(\beta/\mu)$, if $\text{CRSK}(((T, U), \mu)) = ((P, Q), \lambda)$, then $\text{wt}(T) = \text{wt}(P)$ and $\text{wt}(U) = \text{wt}(Q)$.

Proof. For μ, T, and U as in Theorem 5.2, we will construct $\text{CRSK}(((T, U), \mu))$ as follows:

Algorithm 5.3 (Cylindric RSK)**.**

Function CylindricRSK(((tableau T, tableau U), partition μ)): ▷ μ must be the inner shape of both T and U.

1: $P := T$.
2: $\alpha :=$ outer shape of T.
3: $Q :=$ empty tableau with shape α/α.
4: $i :=$ smallest entry that is present in any of the boxes in U.
5: **while** true **do**:
6: $S :=$ set of all boxes that, in U, contain i.
7: $(P, S') :=$ FullMulti$((P, S))$. ▷ That is, P gets assigned to the tableau returned by Algorithm 3.12 and S' gets assigned to the returned set of boxes.
8: Put i into Q in all boxes in S'; expand the outer shape of Q to include all boxes in S'.
9: **if** i is less than the largest entry that is present in any of the boxes in U **then**:
10: $i :=$ smallest entry greater than i that is present in any of the boxes in U.
11: **else**:
12: Exit the loop.
13: **end if**.
14: **end while**.
15: $\lambda :=$ outer shape of P.
16: **return** $((P, Q), \lambda)$.

Example 5.4. Suppose that we start with the following T and U:

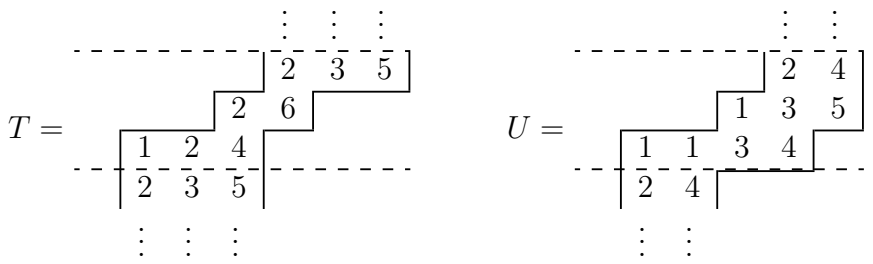

Setting P equal to T and Q equal to the (empty) tableau with shape α/α, we have:

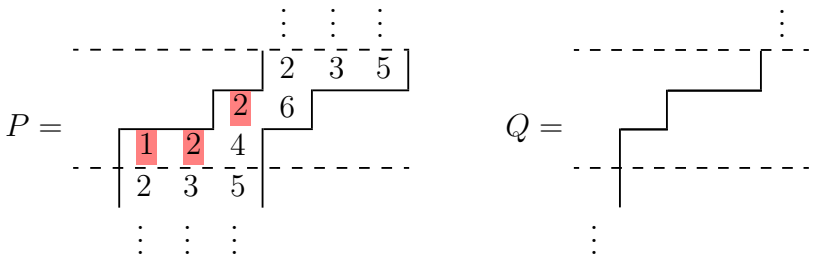

First, $i = 1$. We let S be the set of boxes that, in U, contain 1 (shown in red in the diagram of P above). We then perform Algorithm 3.12 on (P, S), letting P be equal to the return tableau and letting S' be the returned set of boxes. We put 1 in Q in all boxes in S'. The result is shown below. Boxes that are new to P are shown in green; boxes that are outside (to the right of) the inner shape of P, but will be inside (to the left of) the inner shape of P after the next call of Algorithm 3.12, are shown in red.

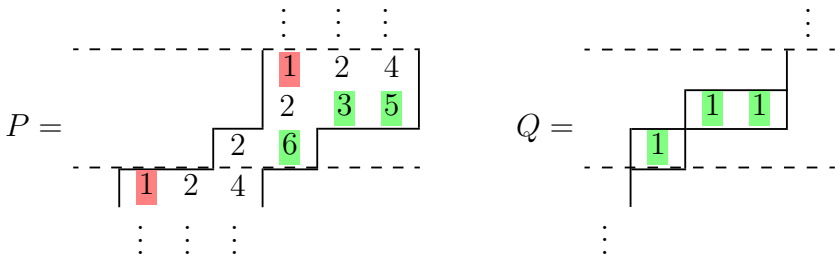

This process continues as follows.

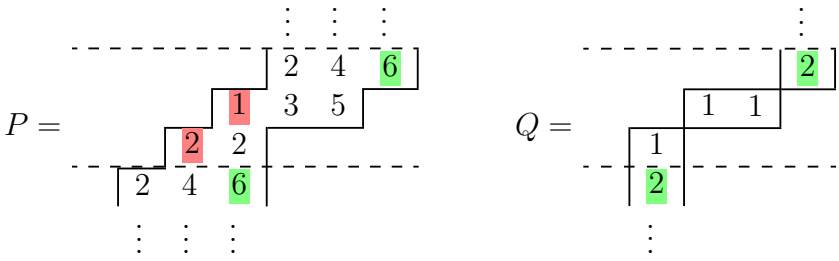

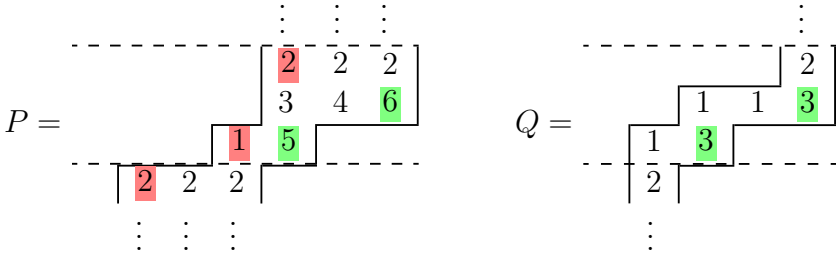

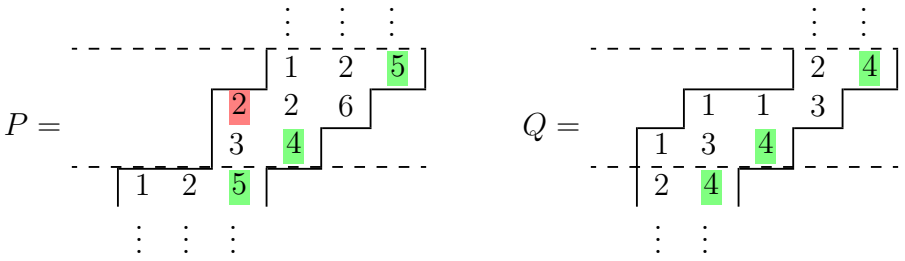

Finally, we have $i = 5$. Since 5 is the largest entry in U, these are our final values of P and Q, and they are returned (in addition to their common outer shape) by Algorithm 5.3.

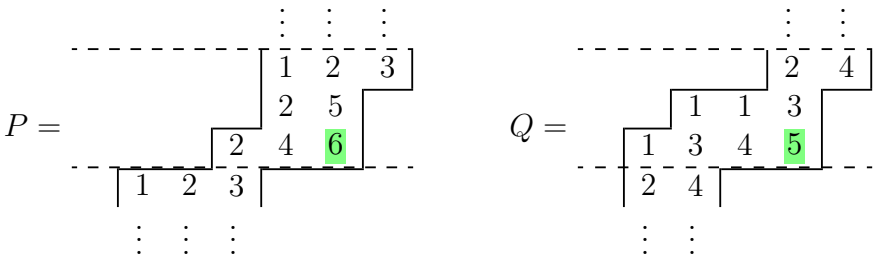

Example 5.5. The following example is shown to demonstrate what happens when, at some time during the execution of Algorithm 5.3 (specifically its subroutine, Algorithm 3.12), the

inner shape of P is expanded to include boxes that are not present in P (this is the "Else" in Algorithm 3.12 on line 15).

$T =$ 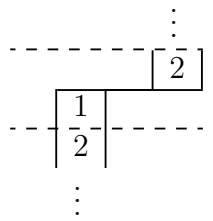

$U =$ 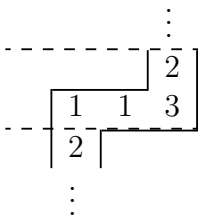

We have

$P =$ 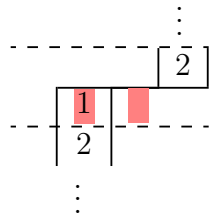

$Q =$ 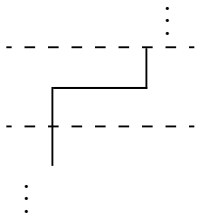

Here, the inner shape of P expands to include a box that is not in P. This means that, despite the fact that two boxes are added to P's inner shape, only one entry (which is $(1,1)$) is added to $\mathfrak{q}$ in the subroutine. The 1 takes the place of the 2 in row 0, putting $(2,0)$ into $\mathfrak{q}$. The 2 is placed in row 1 (which is empty) and gets placed in the first box to the right of P's new inner shape (the box just to the right of the right red box above). However, 1 is added in Q into both boxes that were added to the inner shape of P during the subroutine. Thus, we have:

$P =$ 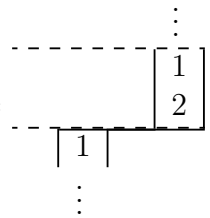

$Q =$ 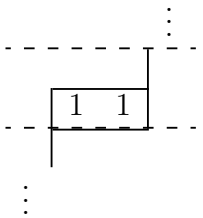

The reader may check that the following is the end result of the algorithm:

$P =$ 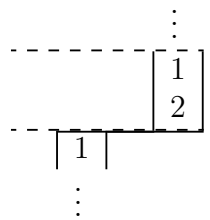

$Q =$ 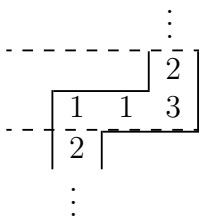

The reader should note that $P \neq T$ and that $Q \neq U$. The boxes in P are shifted down and to the right by one unit from the boxes in T; the boxes in Q are shifter to the right by one unit from the boxes in U.

For λ, P, and Q as in Theorem 5.2, we will construct $\mathrm{CRSK}^{-1}(((P,Q),\lambda))$ as follows:

Algorithm 5.6 (Inverse Cylindric RSK)**.**

Function InverseCylindricRSK(((tableau P, tableau Q), partition λ)): ▷ λ must be the outer shape of both P and Q.

$T := P$.
$\beta :=$ inner shape of P.
$U :=$ empty tableau with shape β/β.
$i :=$ largest entry that is present in any of the boxes in Q.
while true **do**:
    $S :=$ set of all boxes that, in Q, contain i.
    $(T, S') :=$ ReverseFullMulti$((T, S))$. ▷ That is, T gets assigned to the tableau returned by Algorithm 4.7 and S' gets assigned to the returned set of boxes.
    Put i into U in all boxes in S'; shrink the inner shape of U to exclude all boxes in S'.
    if i is greater than the smallest entry that is present in any of the boxes in Q **then**:
        $i :=$ largest entry less than i that is present in any of the boxes in Q.
    else:
        Exit the loop.
    end if.
end while.
$\mu :=$ inner shape of T.
return $((T, U), \mu)$.

For an example of our inverse cylindric RSK algorithm, the reader may look at Example 5.4, starting with the last two tableaux drawn as P and Q, respectively. The pairs of tableaux obtained after each call of the subroutine (Algorithm 4.7) and update of U are the same as the pairs of tableaux in Example 5.4, except in reverse order. The algorithm ends with T and U equalling the T and U given at the beginning of Example 5.4, respectively.

Now we will prove that Algorithm 5.3 and Algorithm 5.6 are valid algorithms for $\mathrm{CRSK}((T, U), \mu)$ and $\mathrm{CRSK}^{-1}((P, Q), \lambda)$, respectively.

Given any $\alpha, \beta, \mu \in$ Cylpar such that $\mu \subseteq \alpha$ and $\mu \subseteq \beta$, as well as any $T \in \mathrm{SSCT}(\alpha/\mu)$ and $U \in \mathrm{SSCT}(\beta/\mu)$, let $\mathrm{CylindricRSK}(((T, U), \mu)) = ((P, Q), \lambda)$. Based on our theorem statement, in order to show that Algorithm 5.3 is a valid algorithm for the CRSK function, we need to show that:

- Every time that Algorithm 3.12 is called from Algorithm 5.3, all preconditions for Algorithm 3.12 are met.
- Algorithm 5.3 terminates;
- $\lambda \in$ Cylpar;
- $\alpha \subseteq \lambda$ and $\beta \subseteq \lambda$;
- $P \in \mathrm{SSCT}(\lambda/\beta)$ and $Q \in \mathrm{SSCT}(\lambda/\alpha)$; and
- $\mathrm{wt}(T) = \mathrm{wt}(P)$ and $\mathrm{wt}(U) = \mathrm{wt}(Q)$.

Given any $\alpha, \beta, \lambda \in$ Cylpar such that $\alpha \subseteq \lambda$ and $\beta \subseteq \lambda$, as well as any $P \in \mathrm{SSCT}(\lambda/\beta)$ and $Q \in \mathrm{SSCT}(\lambda/\alpha)$, let $\mathrm{InverseCylindricRSK}(((P, Q), \lambda)) = ((T, U), \mu)$. Based on our theorem

statement, in order to show that Algorithm 5.6 is a valid algorithm for the inverse of the CRSK function, we need to show that:

- Every time that Algorithm 4.7 is called from Algorithm 5.6, all preconditions for Algorithm 4.7 are met.
- Algorithm 5.6 terminates;
- $\mu \in \text{Cylpar}$;
- $\mu \subseteq \alpha$ and $\mu \subseteq \beta$;
- $T \in \text{SSCT}(\alpha/\mu)$ and $U \in \text{SSCT}(\beta/\mu)$; and
- $\text{wt}(T) = \text{wt}(P)$ and $\text{wt}(U) = \text{wt}(Q)$.

Finally, we need to show that Algorithm 5.6 does indeed invert Algorithm 5.3. That is, we need to show that, for a valid input $((T, U), \mu)$ into Algorithm 5.3,

$$\text{InverseCylindricRSK}(\text{CylindricRSK}(((T, U), \mu))) = ((T, U), \mu),$$

and that, for a valid input $((P, Q), \lambda)$ into Algorithm 5.6,

$$\text{CylindricRSK}(\text{InverseCylindricRSK}(((P, Q), \lambda))) = ((P, Q), \lambda).$$

We will prove these facts below.

Remark 5.7. *Every time that Algorithm 3.12 is called from Algorithm 5.3, all preconditions for Algorithm 3.12 are met.*

Proof. We will proceed by induction. We know that the original P (which equals T) is a valid semistandard tableau. We also know that μ plus the boxes in U that contain the smallest entry of U (call it m) is a valid partition (because otherwise there would be a box in U containing m that would be below or to the right of a box containing an entry less than m). Furthermore, any set of boxes that, in a tableau, contain the same entry, forms a horizontal strip (since otherwise there would be an entry below an entry of the same value in the tableau). Thus, the preconditions for Algorithm 3.12 are satisfied the first time that the subroutine is called.

Suppose that, up to but not including the call of Algorithm 3.12 when $i = j$ (for some $j > m$), the preconditions of the algorithm are satisfied. We know, then, that P is a valid semistandard tableau at the time of the call when $i = j$ (because Algorithm 3.12 returns valid semistandard tableaux). We know that S forms a horizontal strip of boxes (by the same reasoning as above).

Let μ_j be the inner shape of P when $i = j$ but before Algorithm 3.12 is called in that iteration of the loop. We know that μ_j is a valid partition. Suppose, for contradiction, that adding to μ_j the boxes in S will result in an invalid partition. Let μ'_j be the result when these boxes are added (by our assumption, μ'_j is not a valid partition). Then one of the boxes in S (call it B) has a box above it or to its left (call it C) that is neither in μ_j, nor in S.

The boxes that are not in U fall into two categories: boxes in μ and boxes not in β. We know that any box above or to the left of B that is not in U falls into the first category.

Thus, if C is not in U, $C \in \mu$, and $\mu \subseteq \mu_j$. It follows that $C \in \mu_j$, a contradiction. If C is in U but is not in μ_j, then, in U, C contains an entry that is greater than or equal to j. We know that C cannot contain j because then it would have been in S. This means that C contains an entry greater than j in U. This is a contradiction, as B contains j and C is above or to the left of B. Therefore, adding S to μ_j results in a valid partition. This means that all preconditions for Algorithm 3.12 are satisfied.

Having completed our induction, we have shown that all preconditions for Algorithm 3.12 are met every time that the algorithm is called from Algorithm 5.3. □

Remark 5.8. *Algorithm 5.3 terminates.*

Proof. By Remark 3.14, we know that the subroutine called from Algorithm 5.3 always terminates. The loop in Algorithm 5.3 terminates as well, as Q has a finite number of boxes (since λ/μ has a finite number of boxes for all $\mu \subseteq \lambda$) and therefore has a finite number of distinct entries (values that i can take on). Therefore, Algorithm 5.3 always terminates. □

Remark 5.9. $\lambda \in \text{Cylpar}$.

Proof. Since the subroutine (Algorithm 3.12) always returns a valid tableau, the final tableau returned by the subroutine (which is the final value of P) is bounded by a valid outer shape. This is the partition λ that is returned by the algorithm. □

Remark 5.10. $\alpha \subseteq \lambda$ *and* $\beta \subseteq \lambda$.

Proof. Suppose, for contradiction, that $\alpha \not\subseteq \lambda$. Then there is a box that is in α, but is not in λ. Since at the beginning of the algorithm, α is the outer shape of P, and at the end, λ is the outer shape of P, it follows that the outer shape of P lost at some point during the algorithm. P is only modified in the subroutine, but in the subroutine the outer shape of the tableau can only expand. This is a contradiction. Therefore, $\alpha \subseteq \lambda$.

Suppose, for contradiction, that $\beta \not\subseteq \lambda$. Then P would not be a valid tableau. This is a contradiction, as P is returned by Algorithm 3.12, which always returns valid tableaux. Therefore, $\beta \subseteq \lambda$. □

Remark 5.11. $P \in \text{SSCT}(\lambda/\beta)$ *and* $Q \in \text{SSCT}(\lambda/\alpha)$.

Proof. The outer shape of P is λ by definition. The inner shape of P is μ, plus the boxes of U (i.e. the boxes in β/μ). Therefore, the inner shape of P is β. P is a semistandard tableau, as it is returned by Algorithm 3.12, which always returns semistandard tableaux. Therefore, $P \in \text{SSCT}(\lambda/\beta)$.

The boxes of Q are the boxes that are in P but not in T (since they are the boxes that were added to P after every call of Algorithm 3.12). Since $T \in \text{SSCT}(\alpha/\mu)$ and $P \in \text{SSCT}(\lambda/\beta)$, Q consists of all boxes that are in λ but not in α — that is, Q's inner shape is α and its outer shape is λ.

Suppose, for contradiction, that Q is not a semistandard tableau. Then there is a box B in Q that contains an entry (call it b) that is below a box that contains b, or is below or to the right of a box that contains an entry greater than b.

The first case would imply that S' does not form a horizontal strip. This is a contradiction, by Corollary 3.30. In the second case, let C be a box that is above or to the left of B that

contains an entry greater than b. Consider the tableau that is returned by Algorithm 3.12 during the iteration of the loop of Algorithm 5.3 when $i = b$; call this tableau R. B is in S' (and is therefore in the outer shape of R), but C gets added to the outer shape of P on a later call of the loop, and is therefore not in the outer shape of R. This means that the outer shape of R is not a valid partition. This is a contradiction. Therefore, Q is semistandard. It follows that $Q \in \mathrm{SSCT}(\lambda/\alpha)$. □

Remark 5.12. $\mathrm{wt}(T) = \mathrm{wt}(P)$ *and* $\mathrm{wt}(U) = \mathrm{wt}(Q)$.

Proof. Algorithm 3.12 clearly preserves the weight of a tableau. It follows that $\mathrm{wt}(T) = \mathrm{wt}(P)$. In any given iteration of the loop in Algorithm 5.3, $|S|$ is the number of times that i appears in U, and $|S'|$ is the number of times that i appears in Q. Clearly, $|S| = |S'|$. Therefore, $\mathrm{wt}(U) = \mathrm{wt}(Q)$. □

We have proven all necessary facts about Algorithm 5.3. The proofs of the facts outlined above about Algorithm 5.6 have entirely analogous proofs, and, for this reason, are not provided in this paper.

Next, we show that our algorithms are indeed inverses of each other.

Remark 5.13. $\mathrm{InverseCylindricRSK}(\mathrm{CylindricRSK}(((T, U), \mu))) = ((T, U), \mu)$

Proof. Let $((P, Q), \lambda) = \mathrm{CylindricRSK}(((T, U), \mu))$. We first note that $((P, Q), \lambda)$ satisfies the input preconditions for Algorithm 5.6: P and Q are both valid semistandard tableaux with outer shape λ (Remark 5.11).

Define P_j and Q_j to be the values of P and Q, respectively, during the execution of Algorithm 5.3 on $((T, U), \mu)$ while $i = j$, but *before* Algorithm 3.12 is called inside the loop. Define T_j and U_j to be the values of T and U, respectively, during the execution of Algorithm 5.6 on $((P, Q), \lambda)$ while $i = j$, *after* Algorithm 4.7 is called inside the loop. Define S_j to be the set of boxes that, in U, contain j, and N_j to be the set of boxes that, in Q, contain j.

Suppose that, for all $j > m$ for some m, $P_j = T_j$ and U_j is the tableau of boxes that, in U, contain entries greater than or equal to j, with those boxes holding the same entries as they do in U. We will show that $P_m = T_m$ and that U_m is the tableau of boxes that, in U, contain entries greater than or equal to m, with those boxes holding the same entries as they do in U.

If m is the largest value that i takes on (that is, it is the largest entry in any box in U), let $V = P$ and $W = Q$. Otherwise, let M be the smallest entry in U that is greater than m (that is, it is the value that i takes on after m in Algorithm 5.3, and it is the value that i takes on before m in Algorithm 5.6); let $V = T_M$ and $W = U_M$. We know that $\mathrm{FullMulti}(P_m, S_m) = V$, because we know that the tableau returned by Algorithm 3.12 is $P_M = T_M = V$ (or is V if m is the largest value that i takes on). Since N_m is the set of boxes that, in Q, contain m, it follows that $\mathrm{FullMulti}((P_m, S_m)) = (V, N_m)$. V is the tableau from which, during Algorithm 5.6, the boxes that, in Q, contain m are removed (in the iteration of the loop following the one in which V is defined). Thus, $\mathrm{ReverseFullMulti}((V, N_m)) = (T_m, X)$ for some set of boxes X. By Theorem 4.24, it follows that $T_m = P_m$ and that $X = S_m$. From the latter fact it follows that m is placed in U_m in the boxes of S_m. By definition, S_m is the set of boxes that, in U, contain m. Since W is the tableau of boxes that, in U, contains entries *greater* than m, with those boxes holding the same entries as they do in U (by inductive hypothesis

or trivially in the case that m is the largest value that i takes on), it follows that U_m is the tableau of boxes that, in U, contains entries greater than or equal to m, with those boxes holding the same entries as they do in U. Our induction being complete, we have proven that $P_t = T_t$ and $U_t = U$, where t is the smallest value that i takes on. Since T_t and U_t are the final values that T and U take on in Algorithm 5.6, $P_t = T$ in Algorithm 5.3, and μ is the inner shape of T, it follows that $((T, U), \mu)$ is returned by Algorithm 5.6, as desired. □

The proof that CylindricRSK(InverseCylindricRSK(((P, Q), λ))) = ((P, Q), λ) is entirely analogous and, for this reason, is not provided in this paper.

Having proven the validity of the outputs of Algorithms 5.3 and 5.6, the outputs' compliance with the theorem statement, and the fact that the algorithms are inverses of each other, we have shown the existence of a cylindric RSK correspondence.[9] □

5.2. Consequences of the Cylindric RSK Correspondence

Just as the RSK correspondence can be used in order to prove many results about regular tableaux, the cylindric RSK correspondence can be used in order to prove results about cylindric tableaux. Some such results are shown below.

Theorem 5.14 (Cylindric Cauchy Identity)**.** *Given two cylindric partitions α and β, we have*

$$\sum_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha; \mu \subseteq \beta}} s_{\alpha/\mu}(\mathbf{x}) s_{\beta/\mu}(\mathbf{y}) = \sum_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda; \beta \subseteq \lambda}} s_{\lambda/\beta}(\mathbf{x}) s_{\lambda/\alpha}(\mathbf{y}). \tag{4}$$

where $\mathbf{x}$ *and* $\mathbf{y}$ *are variable sets.*

Proof. It is easy to see that

$$\sum_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha; \mu \subseteq \beta}} s_{\alpha/\mu}(\mathbf{x}) s_{\beta/\mu}(\mathbf{y}) = \sum_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha; \mu \subseteq \beta; \\ T \in \text{SSCT}(\alpha/\mu); \\ U \in \text{SSCT}(\beta/\mu)}} \mathbf{x}^{\text{wt}(T)} \mathbf{y}^{\text{wt}(U)} \tag{5}$$

and that

$$\sum_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda; \beta \subseteq \lambda}} s_{\lambda/\beta}(\mathbf{x}) s_{\lambda/\alpha}(\mathbf{y}) = \sum_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda; \beta \subseteq \lambda; \\ P \in \text{SSCT}(\lambda/\beta); \\ Q \in \text{SSCT}(\lambda/\alpha)}} \mathbf{x}^{\text{wt}(P)} \mathbf{y}^{\text{wt}(Q)}. \tag{6}$$

The cylindric RSK correspondence establishes a bijection between such pairs (T, U) and (P, Q) as in equations 5 and 6. Furthermore, the bijection takes a pair (T, U) to a pair (P, Q) such that $\text{wt}(T) = \text{wt}(P)$ and $\text{wt}(U) = \text{wt}(Q)$, meaning that $\mathbf{x}^{\text{wt}(T)} = \mathbf{x}^{\text{wt}(P)}$ and

[9] Having proven the validity of CylindricRSK and InverseCylindricRSK, we can now refer to them as they are referred to in the statement of Theorem 5.2: CRSK and CRSK^{-1}, respectively.

$\mathbf{y}^{\mathrm{wt}(U)} = \mathbf{y}^{\mathrm{wt}(Q)}$. Therefore,

$$\sum_{\substack{\mu\in\mathrm{Cylpar};\\ \mu\subseteq\alpha;\mu\subseteq\beta;\\ T\in\mathrm{SSCT}(\alpha/\mu);\\ U\in\mathrm{SSCT}(\beta/\mu)}} \mathbf{x}^{\mathrm{wt}(T)}\mathbf{y}^{\mathrm{wt}(U)} = \sum_{\substack{\lambda\in\mathrm{Cylpar};\\ \alpha\subseteq\lambda;\beta\subseteq\lambda;\\ P\in\mathrm{SSCT}(\lambda/\beta);\\ Q\in\mathrm{SSCT}(\lambda/\alpha)}} \mathbf{x}^{\mathrm{wt}(P)}\mathbf{y}^{\mathrm{wt}(Q)}.$$

Our theorem statement is therefore proven, by the transitive property of equality. □

Definition 5.15. A *standard cylindric tableau* is a semistandard cylindric tableau whose alphabet is $1, 2, \ldots, m$ for some nonnegative integer m, such that each entry of the tableau's alphabet appears in exactly one box of the tableau. An example of a standard cylindric tableau is shown below.

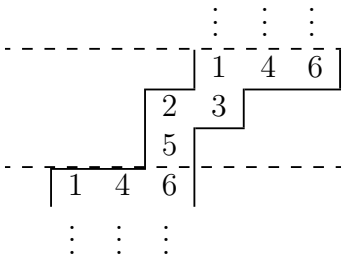

Definition 5.16. Given two cylindric partitions λ and μ, $f_{\lambda/\mu}$ will denote the number of standard tableaux with shape λ/μ.

Corollary 5.17. *Let α and β be two cylindric partitions and m be a nonnegative integer. Let M be the set of all partitions μ such that $\mu \subseteq \alpha$, $\mu \subseteq \beta$, and α/μ contains m distinct boxes. Let Λ be the set of all partitions λ such that $\alpha \subseteq \lambda$, $\beta \subseteq \lambda$, and λ/β contains m distinct boxes. Then*

$$\sum_{\mu\in M} f_{\alpha/\mu} f_{\beta/\mu} = \sum_{\lambda\in\Lambda} f_{\lambda/\alpha} f_{\lambda/\beta}. \tag{7}$$

Proof. Let a be the number of boxes that are in α, but not in β; let b be the number of boxes that are in β, but not in α. Then the number of boxes in β/μ is $m-a+b$. Let $p = m-a+b$. Similarly, the number of boxes in λ/α is p.

For any standard tableau T with shape α/μ, we have $\mathbf{x}^{\mathrm{wt}(T)} = x_1x_2\ldots x_m$. Similarly, for any standard tableau U with shape β/μ, we have $\mathbf{y}^{\mathrm{wt}(U)} = y_1y_2\ldots y_p$.

The left hand side of (7) is the number of pairs of standard tableaux (T, U) with a common inner shape and outer shapes of α and β, respectively. It follows that this is the coefficient of $x_1x_2\ldots x_m \cdot y_1y_2\ldots y_p$ on the left hand side of (4).

Similarly, we have that the right hand side of (7) is the coefficient of $x_1x_2\ldots x_m \cdot y_1y_2\ldots y_p$ on the right hand side of (4). Therefore, by Theorem 5.14, the left hand side and the right hand side of (7) have the same value, as desired. □

5.3. The Symmetry Property of CRSK

We now prove an important property of CRSK, from which we can then derive more identities.

Theorem 5.18. *Given T, U, and μ that satisfy the preconditions of* $\mathrm{CRSK}(((T,U),\mu))$*, let* $\mathrm{CRSK}(((T,U),\mu)) = ((P,Q),\lambda)$*. Then* $\mathrm{CRSK}(((U,T),\mu)) = ((Q,P),\lambda)$*.*

Proof. We first note that it is clear that if the preconditions of $\mathrm{CRSK}(((T,U),\mu))$ are met, then the preconditions of $\mathrm{CRSK}(((U,T),\mu))$ are also met.

We will prove this theorem T and U with alphabets $(1,2,\ldots,m_T)$ and $(1,2,\ldots,m_U)$, respectively; the result for tableaux with other alphabets follows directly by analogy. We will also assume, without loss of generality, that all elements of $\{1,2,\ldots,m_T\}$ appear in T and that all elements of $\{1,2,\ldots,m_U\}$ appear in U (this can be assumed because otherwise the entries in T or U can be "compressed" and m_T or m_U reduced).

Define T_0 to be T and T_j, for $1 \leq j \leq m_U$, to be the value of P after line 7 of Algorithm 5.3 is performed (with input $((T,U),\mu)$), when $i = j$. Define U_0 to be U and U_j, for $1 \leq j \leq m_T$, to be the value of P after line 7 of Algorithm 5.3 is performed (with input $((U,T),\mu)$), when $i = j$.

Define $\lambda_{i,j}$, for $0 \leq i \leq m_U$ and $0 \leq j \leq m_T$, to be the outer shape of the tableau with the same inner shape as T_i and containing exactly the boxes that, in T_i, contain entries that are less than or equal to j. (That is, $\lambda_{i,j}$ is the right bound on the entries 1 through j in T_i.) Define $\nu_{i,j}$, for $0 \leq i \leq m_T$ and $0 \leq j \leq m_U$, to be the outer shape of the tableau with the same inner shape as U_i and containing exactly the boxes that, in U_i, contain entries that are less than or equal to j.

We will perform induction on two statements, for all $1 \leq c \leq m_U$ and $1 \leq d \leq m_T$:

(1) Let B be a box. Then d is bumped from B during the formation of T_c if and only if c is bumped from B during the formation of U_d; and

(2) $\lambda_{c,d} = \nu_{d,c}$.

Suppose that $d = 0$. Then $\lambda_{c,d}$ is the inner shape of T_c. Based on how CRSK works, we know that $\lambda_{c,0}$ is the outer shape of the tableau with inner shape that of U, containing exactly the boxes that, in U, contain entries that are less than or equal to c. It is easy to see that this equals $\nu_{0,c} = \nu_{d,c}$. Similarly, if $c = 0$, then $\lambda_{c,d} = \nu_{d,c}$. These two facts serve as our base cases (as it will turn out, base cases for statement (1) are not necessary).

Suppose that statements (1) and (2) are true for all $d < j$, for some $1 \leq j \leq m_T$. We will show that they are true for $d = j$ as well.

Suppose that statements (1) and (2) are true for $d = j$ and all $c < i$, for some $1 \leq i \leq m_U$. We will show that they are true for $d = j$ and $c = i$ as well.

We will first prove statement (1) for i and j — that is, that, for any box B, j is bumped from B during the formation of T_i if and only if i is bumped from B during the formation of U_j.

Choose a particular box B. We will first prove that if j is bumped from B during the formation of T_i, then i is bumped from B during the formation of U_j.

Case 1: In T, B contains j. Since j is bumped from B during the formation of T_i (and is replaced by an entry that is less than j), we know that $B \in \lambda_{i,j-1}/\lambda_{i-1,j-1}$. It follows from our inductive hypothesis for statement (2) that $B \in \nu_{j-1,i}/\nu_{j-1,i-1}$, which means that, in U_{j-1}, B contains i. By how CRSK works, we know that an entry is bumped out of B during the formation of U_j (since T contains j in B). Thus, i is bumped from B during the formation of U_j, as desired.

Case 2: In T, B does not contain j. B contains j in T_{i-1}, so $B \in \lambda_{i-1,j}/\lambda_{i-1,j-1}$. By our inductive hypothesis on statement (2), it follows that $B \in \nu_{j,i-1}/\nu_{j-1,i-1}$. Thus, an entry that is greater than or equal to i is bumped out of B during the formation of U_j. (From Corollary 3.26, we know that only one entry is bumped out of B during the formation of U_j.) It remains to show that this entry is indeed i.

Suppose that in U, B contains i. Then in U_{j-1}, B contains an entry that is less than or equal to i. However, an entry that is greater than or equal to i is bumped out of B during the formation of U_j. Thus, i must be the entry that is bumped out of B during the formation of U_j, as desired.

Now suppose that in U, B does not contain i. This means that, during the formation of T_i, B is not one of the boxes that is originally removed in the subroutine (that is, $B \notin S$). That is, j is bumped out of B by an entry j'' during the formation of T_i. This means that $B \in \lambda_{i,j''}/\lambda_{i-1,j''}$. By our inductive hypothesis on statement (2), this means that $B \in \nu_{j'',i}/\nu_{j'',i-1}$ — that is, that B contains i in $U_{j''}$. For all $j'' < x < j$, x is not bumped from B during the formation of T_i; it follows from our inductive hypothesis on statement (1) that for all $j'' < x < j$, i is not bumped from B during the formation of U_x. Therefore, i is in B in U_{j-1}. We conclude that i is the entry that is bumped out of B during the formation of U_j in this case as well.

Note that we used our inductive hypothesis on statement (1) once, and we used it for $c = i$ and $d = x$ for $j'' < x < j$ (that is, for defined values of c and d); for this reason, we did not need a base case for this proof.

The converse of the statement that we have proven follows by symmetry. This is not immediately trivial, because our induction is not symmetric (we induct on d and, for each d, induct on c); however, the reader may notice that, in all cases, the inductive hypothesis was used only for $c \leq i$ and $d \leq j$. Thus, the converse of the statement does indeed follow by symmetry.

We have shown statement (1) to be true for $c = i$ and $d = j$. Now we show that statement (2) is true for $c = i$ and $d = j$.

Consider the process of formation of T_i and consider a particular row r. By Corollary 3.27, we know that all j's that are bumped into r are bumped into r after all entries less than j are bumped into r. It follows that all j's that are bumped into r are inserted just to the right of the rightmost of the following:

- The rightmost entry less-than-or-equal-to j that is in T_{i-1}; or
- The rightmost entry less than j that is in T_i (in the case that, during the formation of T_i, entries less than j bump out all j's that were in T_{i-1}).

That is, the j's are placed just to the right of the rightmost of $\lambda_{i-1,j}$ and $\lambda_{i,j-1}$ in row r.

Similarly, during the formation of U_j, the i's are placed in row r just to the right of the rightmost of $\nu_{j-1,i}$ and $\nu_{j,i-1}$. By our inductive hypothesis, $\lambda_{i-1,j} = \nu_{j,i-1}$ and $\lambda_{i,j-1} = \nu_{j-1,i}$.

Since, for any box B in row $r-1$, j is bumped from B during the formation of T_i if and only if i is bumped from B during the formation of U_j, it follows that the same number of j's are inserted into r during the formation of T_i as the number of i's that are inserted into r during the formation of U_j. Since the leftmost of these insertion locations is the same (see the previous paragraph), we conclude that the j's inserted during the formation of T_i are

inserted into the same boxes as the i's inserted during the formation of U_j. It follows that $\lambda_{i,j} = \nu_{j,i}$, as desired.

We have completed our induction on i and have shown that statements (1) and (2) are true for $d = j$ and all values of c.

We have completed our induction on j and have shown that statements (1) and (2) are true for all values d and all values of c.

Let $\text{CRSK}(((T, U), \mu)) = ((P, Q), \lambda)$ and $\text{CRSK}(((U, T), \mu)) = ((Q', P'), \lambda')$. We will show that $P = P'$ (it will follow by analogy that $Q = Q'$, and it will follow that $\lambda = \lambda'$).

Define γ_i, for $1 \leq i \leq m_T$, to be the outer shape of the tableau with the same inner shape as P and containing exactly the boxes that, in P, contain entries that are less than or equal to i; define γ_0 to be the inner shape of P. Define γ'_i analogously for P'. It suffices to show that, for all i, $\gamma_i = \gamma'_i$.

γ_0 and γ'_0 are both the outer shape of U, so clearly $\gamma_0 = \gamma'_0$. For any $1 \leq i \leq m_T$, $\gamma_i = \lambda_{m_U,i}$ (since $P = T_{m_U}$). By how CRSK works, we know that, for any $1 \leq i \leq m_T$, γ'_i is the outer shape of U_i. Thus, $\gamma'_i = \nu_{i,m_U}$. By statement (2), we know that $\lambda_{m_U,i} = \nu_{i,m_U}$. Therefore, $\gamma_i = \gamma'_i$, as desired.

Therefore, for any two tableaux T and U and partition μ that satisfy the preconditions of $\text{CRSK}(((T, U), \mu))$ and $\text{CRSK}((U, T), \mu)$, if $\text{CRSK}(((T, U), \mu)) = ((P, Q), \lambda)$, then $\text{CRSK}(((U, T), \mu)) = ((Q, P), \lambda)$, as desired. □

Corollary 5.19. *Given two tableaux P and Q and a partition λ that satisfy the preconditions of $\text{CRSK}^{-1}(((P, Q), \lambda))$ and $\text{CRSK}^{-1}((Q, P), \lambda)$, let $\text{CRSK}^{-1}(((P, Q), \lambda)) = ((T, U), \mu)$. Then $\text{CRSK}^{-1}(((Q, P), \lambda)) = ((U, T), \mu)$.*

Proof. We have that $\text{CRSK}(((T, U), \mu)) = ((P, Q), \lambda)$. By Theorem 5.18, we have that $\text{CRSK}(((U, T), \mu)) = ((Q, P), \lambda)$. Therefore, $\text{CRSK}^{-1}(((Q, P), \lambda)) = ((U, T), \mu)$, as desired. □

Corollary 5.20. *Given a tableau T with inner shape μ, $\text{CRSK}(((T, T), \mu)) = ((P, P), \lambda)$, for some P and λ.*

Proof. Let $\text{CRSK}(((T, T), \mu)) = ((P, Q), \lambda)$. We are to show that $P = Q$. By Theorem 5.18, we have $\text{CRSK}(((T, T), \mu)) = ((Q, P), \lambda)$. Thus, $P = Q$, as desired. □

Corollary 5.20 shows that CRSK can be viewed as a bijection that takes a tableau with outer shape α to a tableau with inner shape α with the weight. This allows us to obtain some more important results:

Corollary 5.21. *Given any cylindric partition α, we have*

$$\sum_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha}} s_{\alpha/\mu}(\mathbf{x}) = \sum_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda}} s_{\lambda/\alpha}(\mathbf{x}). \tag{8}$$

It is worth mentioning that setting $\beta := \alpha$ and $\mathbf{x} = \mathbf{y}$ in Theorem 5.14 gives us the following related (but different) result:

$$\sum_{\substack{\mu \in \text{Cylpar}; \\ \mu \subseteq \alpha}} s^2_{\alpha/\mu}(\mathbf{x}) = \sum_{\substack{\lambda \in \text{Cylpar}; \\ \alpha \subseteq \lambda}} s^2_{\lambda/\alpha}(\mathbf{x}).$$

We can use (8) in order to obtain the following result, analogous to Corollary 5.17:

Corollary 5.22. *Let α be a cylindric partition and m be a nonnegative integer. Let M be the set of all partitions μ such that $\mu \subseteq \alpha$ and α/μ contains m distinct boxes. Let Λ be the set of all partitions λ such that $\alpha \subseteq \lambda$ and λ/α contains m distinct boxes. Then*

$$\sum_{\mu \in M} f_{\alpha/\mu} = \sum_{\lambda \in \Lambda} f_{\lambda/\alpha}.$$

Relatedly, we can set $\beta := \alpha$ in (7) to obtain:

$$\sum_{\mu \in M} f^2_{\alpha/\mu} = \sum_{\lambda \in \Lambda} f^2_{\lambda/\alpha}.$$

6. A Marble-Game Interpretation of Cylindric Tableaux

Definition 6.1. Consider any partition α. The *marble arrangement* of α (denoted $\text{Arr}(\alpha)$) is constructed as follows. We begin with k people — call them $p_0, p_1, \ldots, p_{k-1}$ — in a circle (k being the vertical period of the cylinder) such that p_i is clockwise from p_{i-1} (note that, just as with rows, p_i refers to $p_{i \pmod k}$, and thus p_0 is clockwise from p_{k-1}). For all i, p_i has $\alpha_{i-1} - \alpha_i$ marbles (α_i is the i'th term in the α's sequence).[10]

Let R be a cylindric tableau with alphabet $\{1, 2, \ldots, t\}$, inner shape μ, and outer shape λ. Define $\text{Arr}_0(R) := \text{Arr}(\mu)$. Next, for $1 \leq j \leq t$, obtain $\text{Arr}_j(R)$ from $\text{Arr}_{j-1}(R)$ as follows: for every row r, let x be the number of times j appears in row r in R. Then p_r in $\text{Arr}_{j-1}(R)$ passes x marbles to p_{r+1}.

For $0 \leq j \leq t$, let $\lambda_j(R)$ be the outer shape of the tableau with inner shape μ and containing exactly the boxes that, in R, contains entries less than or equal to j.

Remark 6.2. *For $0 \leq j \leq t$, $\text{Arr}_j(R) = \text{Arr}(\lambda_j(R))$.*

Proof. We will proceed by induction. We know that $\lambda_0(R) = \mu$, so we have $\text{Arr}_0(R) = \text{Arr}(\lambda_0(R))$. Suppose that $\text{Arr}_j(R) = \text{Arr}(\lambda_j(R))$ for all $0 \leq j < m$. We will show that $\text{Arr}_m(R) = \text{Arr}(\lambda_m(R))$.

Consider any integer r. Say that m appears x_1 times in row r (the projection of plane row r onto the cylinder) and x_2 times in row $r-1$. Then $(\lambda_m(R))_r - (\lambda_{m-1}(R))_r = x_1$ and $(\lambda_m(R))_{r-1} - (\lambda_{m-1}(R))_{r-1} = x_2$. Subtracting these two equations, we have

$$(\lambda_m(R))_{r-1} - (\lambda_m(R))_r = (\lambda_{m-1}(R))_{r-1} - (\lambda_{m-1}(R))_r + x_2 - x_1.$$

Therefore, we have that p_r has $x_2 - x_1$ more marbles in $\text{Arr}(\lambda_m(R))$ than he did in $\text{Arr}(\lambda_{m-1}(R))$.

[10] Clearly, there is a total of $n - k$ marbles ($n - k$ is the horizontal period of the cylinder).

When $\mathrm{Arr}_m(R)$ is obtained from $\mathrm{Arr}_{m-1}(R)$, p_r passes x_1 marbles to p_{r+1} and receives x_2 marbles from p_{r-1}. Thus, we have that p_r has $x_2 - x_1$ more marbles in $\mathrm{Arr}_m(R)$ than he did in $\mathrm{Arr}_{m-1}(R)$. Since this is true for all integers r, and by our inductive hypothesis, we have that $\mathrm{Arr}_m(R) = \mathrm{Arr}(\lambda_m(R))$, as desired. Having completed our induction, we have proven that for $0 \leq j \leq t$, $\mathrm{Arr}_j(R) = \mathrm{Arr}(\lambda_j(R))$. □

Remark 6.3. *For $1 \leq j \leq t$, the formation of $\mathrm{Arr}_j(R)$ never entails a person passing more marbles than he has to this right.*

Proof. The number of marbles that p_r has, being the difference between the length of the row above r and the length of row r in the sub-tableau of R consisting of the entries 1 through $j-1$, describes how many entries j row r can possibly have without the tableau being non-semistandard. □

Definition 6.4. A *turn* is a combination of simultaneous marble passes among the k people, such that each person passes at most as many marbles as he has. A turn is denoted $(a_0, a_1, \ldots, a_{k-1})$, where a_i is the number of marbles that p_i passes to p_{i+1}.

Definition 6.5. A *marble game* of length t is an initial arrangement of $n-k$ marbles among the k people, along with a sequence of t successive turns.

Proposition 6.6. *Given a partition μ and a nonnegative integer t, there is a bijection between marble games of length t with initial arrangement $\mathrm{Arr}(\mu)$ and cylindric tableaux with alphabet $\{1, 2, \ldots, t\}$ and inner shape μ.*

Proof. We have shown a construction of a marble game given a tableau, and, as stated in Remark 6.3, the game thus constructed is a valid one. Given a marble game G, one can obtain a unique tableau that produces G with the process described above. This is because turn j in G describes the horizontal strip of j's that is in its corresponding tableau: from $\lambda_{j-1}(R)$ we obtain $\lambda_j(R)$ by adding to row r the number of j's equal to the number of marbles that p_r gives to p_{r+1}. Because p_r cannot give more marbles than he has to p_{r+1}, the tableau thus produced is indeed a valid semistandard tableau. □

Example 6.7. Let R be the following cylindric tableau, with $t = 6$ (recall that the top row of the tableau is row 0).

$R =$

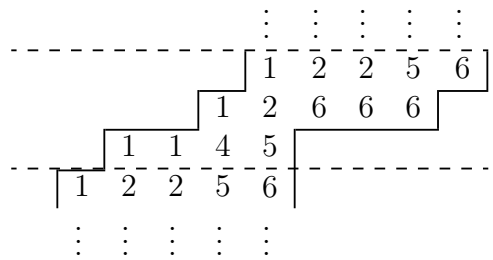

Originally, we have $p_0 = 1$, $p_1 = 1$, and $p_2 = 2$. On the first turn, one 1 is added to rows 0 and 1 and two 1's are added to row 2; thus, after the first turn (represented as $(1, 1, 2)$), we have $p_0 = 2$, $p_1 = 1$, and $p_2 = 1$. The next turn is $(2, 1, 0)$, so after turn 2, we have $p_0 = 0$, $p_1 = 2$, and $p_2 = 2$. There are no 3's in R, so the next turn — $(0, 0, 0)$ — results in $p_0 = 0$,

$p_1 = 2$, and $p_2 = 2$. Turn 4 is $(0, 0, 1)$; turn 5 is $(1, 0, 1)$; and turn 6 is $(1, 3, 0)$. Thus, the marble game corresponding to R has the following sequence of turns:

$$(1, 1, 2), (2, 1, 0), (0, 0, 0), (0, 0, 1), (1, 0, 1), (1, 3, 0).$$

If we know μ, we can use this information to retrace R: the first turn encodes that R has one 1 in rows 0 and 1 and two 1's in row 2; the second turn encodes that R has two 2's in row 0, one 2 in row 1, and no 2's in row 2; and so on.

Note that knowing μ and $\mathrm{Arr}_j(R)$ for $1 \leq j \leq t$ does not uniquely determine R. For example, if $\mathrm{Arr}_1(R) = \mathrm{Arr}(\mu)$, one cannot necessarily determine whether there are no 1's in R, each row has one 1 in R, each row has two 1's in R, etc. Thus, a game, and not just the series of arrangements attained by a game, is necessary in order to describe a tableau.

We can use Proposition 6.6 in order to obtain some results relating the marble game and cylindric tableaux.

Corollary 6.8. *Let μ be a cylindric partition and t be a nonnegative integer. The number of cylindric tableaux with inner shape μ and alphabet $\{1, 2, \ldots, t\}$ equals the number of possible marble games with t turns that begins with the arrangement* $\mathrm{Arr}(\mu)$.

Corollary 6.9. *Let μ be a cylindric partition and t be a nonnegative integer. The number of standard cylindric tableaux with inner shape μ and alphabet $\{1, 2, \ldots, t\}$ equals the number of possible marble games with t turns that begins with the arrangement* $\mathrm{Arr}(\mu)$ *and in which exactly one marble changes hands on every turn.*

Corollary 6.10. *Let λ be a cylindric partition and t be a nonnegative integer. The number of cylindric tableaux with outer shape λ and alphabet $\{1, 2, \ldots, t\}$ equals the number of possible marble games with t turns that ends with the arrangement* $\mathrm{Arr}(\lambda)$.

Note that the number of possible marble games with t turns that ends with the arrangement $\mathrm{Arr}(\lambda)$ is the same as the number of possible marble games with t turns that begins with the arrangement $\mathrm{Arr}(\lambda)$, except that marbles are passed counterclockwise instead of clockwise. This is because this modified marble game retraces the steps of any marble game that ends with $\mathrm{Arr}(\lambda)$.

Corollary 6.11. *Let λ be a cylindric partition and t be a nonnegative integer. The number of standard cylindric tableaux with outer shape λ and alphabet $\{1, 2, \ldots, t\}$ equals the number of possible marble games with t turns that ends with the arrangement* $\mathrm{Arr}(\lambda)$ *and in which exactly one marble changes hands on every turn.*

Corollary 6.12. *Given a cylindric partition α and nonnegative integer t, the number of possible marble games with t turns that starts with the arrangement* $\mathrm{Arr}(\alpha)$ *is equal to the number of possible marble games with t turns that ends with the arrangement* $\mathrm{Arr}(\alpha)$ *(which is in turn equal to the number of possible marble games with t turns that starts with the arrangement* $\mathrm{Arr}(\alpha)$*, but in which marbles are passed counterclockwise). This remains true if we restrict all turns so that exactly one marble changes hands on every turn.*

Proof. If in Corollary 5.21, we let $\mathbf{x} = (x_1, x_2, \ldots, x_t)$ and we let $x_1 = x_2 = \cdots = x_t = 1$, we find that the number of cylindric tableaux with alphabet $\{1, 2, \ldots, t\}$ and inner shape α equals the number of cylindric tableaux with alphabet $\{1, 2, \ldots, t\}$ and outer shape α. The first part of our proposition now follows directly from corollaries 6.8 and 6.10. The second part of our proposition follows analogously from corollaries 5.22, 6.9, and 6.11. □

It is noteworthy that this marble-game construction only works for cylindric tableaux. A similar concept can be defined for regular tableaux, but it would involve an infinite line of people in which the first person can obtain an arbitrary number of marbles — a construction that is not nearly as interesting and much less likely to be of use. The purely combinatorial construction which we have described in this section shows the promise of cylindric tableaux in potential applications outside of tableau theory.

One such application may be in information theory, where marbles are bits and players are computers that are linked in a ring in which communication is only allowed in one direction. In such an interpretation, the numbers in a column represent the times when a particular bit is transferred from one computer to the next. The ease of tracking particular bits with tableau representations of communication in unidirectional rings makes it possible for such representations to be useful in optimizing such communication.

7. Applying Results Concerning Cylindric Tableaux to Skew Tableaux

At the beginning of the paper, we fixed n and k. For this section, we will unfix n and k.

It turns out that results concerning skew tableaux can be proven from analogous results concerning cylindric tableaux. The typical construction for such proofs is, given a skew tableau or shape, to create a cylindric tableau that looks like the skew tableau or shape, but with a very large k (as large as necessary) and an even larger n (so that $n - k$ is as large as necessary). As an example, given the skew shape below on the left, the cylindric shape that would be produced looks like the cylindric shape below on the right.

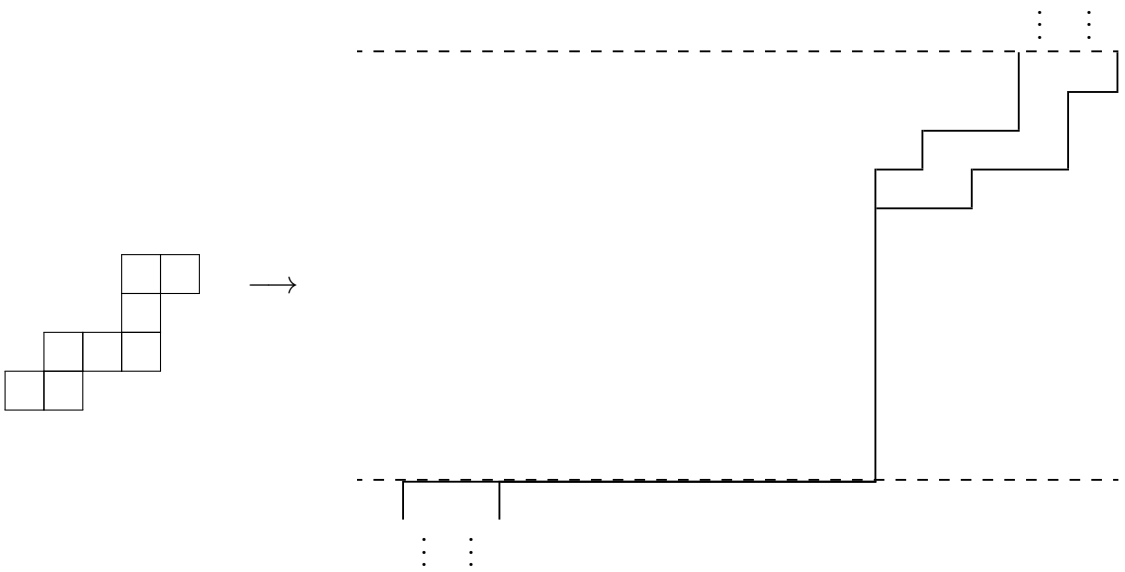

Definition 7.1. If a partition λ is turned into a cylindric partition as shown above, where the vertical period of the cylindric partition is k and the horizontal period of the cylindric partition is $n-k$, we denote the cylindric partition $\mathrm{Cyl}_{k,n}(\lambda)$. (Thus, if above our skew shape is λ/μ, then the cylindric shape shown above is $\mathrm{Cyl}_{k,n}(\lambda)/\mathrm{Cyl}_{k,n}(\mu)$.)

We will illustrate an example of such a proof. The proposition below is not a new result and follows directly from [SagStan, §6, Corollary 6.12], given [Ful, §4.3, Equation (3)]. The purpose of the proposition below is, instead, to demonstrate a proof technique that uses results concerning cylindric tableaux in order to prove results about skew tableaux.

Proposition 7.2. *Given two (non-cylindric) partitions α and β, we have*

$$\sum_{\mu} s_{\alpha/\mu}(\mathbf{x})s_{\beta/\mu}(\mathbf{y}) \cdot \sum_{\gamma} s_{\gamma}(\mathbf{x})s_{\gamma}(\mathbf{y}) = \sum_{\lambda} s_{\lambda/\beta}(\mathbf{x})s_{\lambda/\alpha}(\mathbf{y}).$$

Proof. Given a power series P, let $\mathrm{hg}_d P$ be the degree-d homogeneous component of P (the part of the power series consisting of only degree-d terms). We show that for all d,

$$\mathrm{hg}_d \sum_{\mu} s_{\alpha/\mu}(\mathbf{x})s_{\beta/\mu}(\mathbf{y}) \cdot \sum_{\gamma} s_{\gamma}(\mathbf{x})s_{\gamma}(\mathbf{y}) = \mathrm{hg}_d \sum_{\lambda} s_{\lambda/\beta}(\mathbf{x})s_{\lambda/\alpha}(\mathbf{y}).$$

Consider any particular d. Choose k and n such that k and $n-k$ are both very large compared to d and the dimensions of α and β. Specifically, if $\alpha = (a_1, a_2, \ldots, a_m)$ and $\beta = (b_1, b_2, \ldots, b_p)$, then the following values of k and n should be large enough:

$$k = \max(m, p) + 2d + 1; \qquad n = k + \max(a_1, b_1) + 2d + 1.$$

Lemma 7.3. *Given two partitions α and β, let $\alpha' = \mathrm{Cyl}_{k,n}(\alpha)$ and $\beta' = \mathrm{Cyl}_{k,n}(\beta)$. Then*

$$\mathrm{hg}_d \sum_{\mu} s_{\alpha/\mu}(\mathbf{x})s_{\beta/\mu}(\mathbf{y}) \sum_{\gamma} s_{\gamma}(\mathbf{x})s_{\gamma}(\mathbf{y}) = \mathrm{hg}_d \sum_{\substack{\mu' \in \mathrm{Cylpar}; \\ \mu' \subseteq \alpha'; \mu' \subseteq \beta'}} s_{\alpha'/\mu'}(\mathbf{x})s_{\beta'/\mu'}(\mathbf{y}).$$

(Partitions on the left side above are regular; those on the right side are cylindric.)

Proof. Take any particular μ' that produce terms of degree d on the right hand side of the above equation. Note that this puts a limit on the number of boxes in α'/μ' (and β'/μ'). This means that α'/μ' looks something like this:

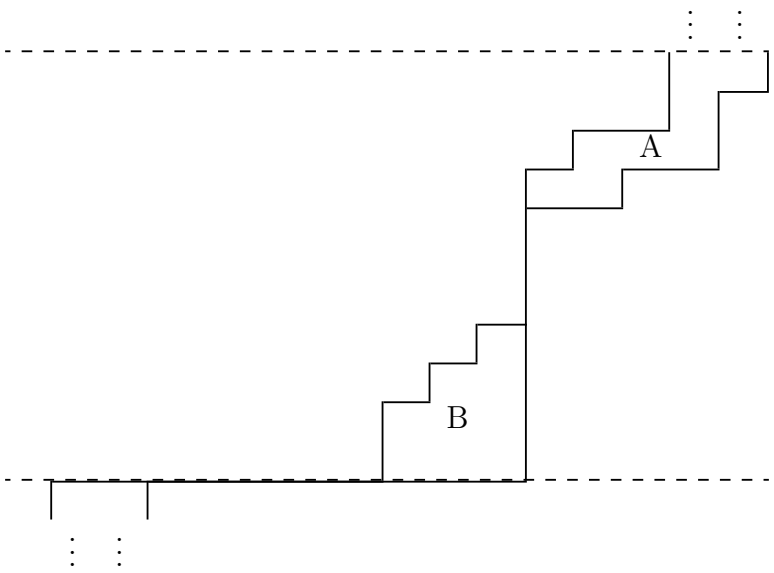

Because k and $n-k$ are large compared to d and the dimensions of α', the regions designated above as A and B are not connected.

Both regions A and B are in α'; neither is in μ'. Let ν' be the partition such that $\mu' \subseteq \nu' \subseteq \alpha'$, B is in ν', and A is entirely outside of ν'. Since A and B are disconnected regions, we have

$$s_{\alpha'/\mu'}(\mathbf{x}) = s_{\alpha'/\nu'}(\mathbf{x}) s_{\nu'/\mu'}(\mathbf{x}).$$

Similarly, we have

$$s_{\beta'/\mu'}(\mathbf{x}) = s_{\beta'/\nu'}(\mathbf{x}) s_{\nu'/\mu'}(\mathbf{x}).$$

Since μ' can be thought of in terms of as α' sans region A, and then sans region B (for small enough sizes of A and B), we have

$$\mathrm{hg}_d \sum_{\substack{\mu' \in \mathrm{Cylpar}; \\ \mu' \subseteq \alpha'; \mu' \subseteq \beta'}} s_{\alpha'/\mu'}(\mathbf{x}) s_{\beta'/\mu'}(\mathbf{y}) = \mathrm{hg}_d \sum_{\nu'} s_{\alpha'/\nu'}(\mathbf{x}) s_{\beta'/\nu'}(\mathbf{y}) \sum_{\mu'} s_{\nu'/\mu'}(\mathbf{x}) s_{\nu'/\mu'}(\mathbf{y}), \qquad (9)$$

where ν' ranges over all cylindric partitions such that $\mathrm{Cyl}_{k,n}(\nu) = \nu'$ for some partition ν such that α/ν (a skew shape) contains d or fewer boxes, and where μ' ranges over all cylindric partitions that are ν', except with a "corner" (such as region B above) — which has at most d boxes — removed from ν'.

For any particular ν' and μ' (satisfying the limitations described in the above paragraph), let $\mathcal{B}$ be the skew shape that looks like the region B between μ' and ν' (so in the above example, $\mathcal{B} = (3,3,3,3)/(2,1)$). Clearly,

$$s_{\nu'/\mu'}(\mathbf{x}) = s_{\mathcal{B}}(\mathbf{x}).$$

Let $\mathcal{B}'$ be the shape that is $\mathcal{B}$ rotated by 180° (because of the shape of the outer shape of $\mathcal{B}$, we have that $\mathcal{B}'$ is a straight (non-skew) partition; in our example, we would have $\mathcal{B}' = (3,3,2,1)$). As it turns out, $s_{\mathcal{B}}(\mathbf{x}) = s_{\mathcal{B}'}(\mathbf{x})$ [GriRei, §2, Exercise 2.22 (b)]. This is because if a skew tableau R with shape $\mathcal{B}$ has content $(a_1, a_2, \ldots, a_m)$, we can create the

tableau R' with shape $\mathcal{B}'$ that has content $(a_m, a_{m-1}, \ldots, a_1)$ by switching all 1's with m's, 2's with $m-1$'s, etc., and by symmetry of Schur polynomials, we have that the number of tableaux with shape $\mathcal{B}'$ and content $(a_m, a_{m-1}, \ldots, a_1)$ equals the number of tableaux with shape $\mathcal{B}'$ and content $(a_1, a_2, \ldots, a_m)$. Therefore, we have that, given ν', for small enough degrees (including all up to d),

$$\sum_{\mu'} s_{\nu'/\mu'}(\mathbf{x}) s_{\nu'/\mu'}(\mathbf{y}) = \sum_{\mathcal{B}'} s_{\mathcal{B}'}(\mathbf{x}) s_{\mathcal{B}'}(\mathbf{y}) = \sum_{\gamma} s_{\gamma}(\mathbf{x}) s_{\gamma}(\mathbf{y}). \tag{10}$$

Combining equations 9 and 10, we have

$$\begin{aligned} \mathrm{hg}_d \sum_{\substack{\mu' \in \mathrm{Cylpar}; \\ \mu' \subseteq \alpha'; \mu' \subseteq \beta'}} s_{\alpha'/\mu'}(\mathbf{x}) s_{\beta'/\mu'}(\mathbf{y}) &= \mathrm{hg}_d \sum_{\nu'} s_{\alpha'/\nu'}(\mathbf{x}) s_{\beta'/\nu'}(\mathbf{y}) \sum_{\gamma} s_{\gamma}(\mathbf{x}) s_{\gamma}(\mathbf{y}) \\ &= \mathrm{hg}_d \sum_{\nu'} s_{\alpha'/\nu'}(\mathbf{x}) s_{\beta'/\nu'}(\mathbf{y}) \sum_{\gamma} s_{\gamma}(\mathbf{x}) s_{\gamma}(\mathbf{y}). \end{aligned}$$

Finally, since $s_{\alpha'/\nu'}(\mathbf{x})$, for every ν', describes the ways to fill the corresponding region A (or its corresponding skew shape), and analogously for β', we have that, for small enough degrees (including all up to d),

$$\sum_{\nu'} s_{\alpha'/\nu'}(\mathbf{x}) s_{\beta'/\nu'}(\mathbf{y}) = \sum_{\mu} s_{\alpha/\mu}(\mathbf{x}) s_{\beta/\mu}(\mathbf{y}).$$

Therefore,

$$\mathrm{hg}_d \sum_{\mu} s_{\alpha/\mu}(\mathbf{x}) s_{\beta/\mu}(\mathbf{y}) \sum_{\gamma} s_{\gamma}(\mathbf{x}) s_{\gamma}(\mathbf{y}) = \mathrm{hg}_d \sum_{\substack{\mu' \in \mathrm{Cylpar}; \\ \mu' \subseteq \alpha'; \mu' \subseteq \beta'}} s_{\alpha'/\mu'}(\mathbf{x}) s_{\beta'/\mu'}(\mathbf{y}),$$

as desired. □

By Theorem 5.14, we have that

$$\mathrm{hg}_d \sum_{\substack{\mu' \in \mathrm{Cylpar}; \\ \mu' \subseteq \alpha'; \mu' \subseteq \beta'}} s_{\alpha'/\mu'}(\mathbf{x}) s_{\beta'/\mu'}(\mathbf{y}) = \mathrm{hg}_d \sum_{\substack{\lambda' \in \mathrm{Cylpar}; \\ \alpha' \subseteq \lambda'; \beta' \subseteq \lambda'}} s_{\lambda'/\beta'}(\mathbf{x}) s_{\lambda'/\alpha'}(\mathbf{y}).$$

Thus, it suffices to show that

$$\mathrm{hg}_d \sum_{\substack{\lambda' \in \mathrm{Cylpar}; \\ \alpha' \subseteq \lambda'; \beta' \subseteq \lambda'}} s_{\lambda'/\beta'}(\mathbf{x}) s_{\lambda'/\alpha'}(\mathbf{y}) = \mathrm{hg}_d \sum_{\lambda} s_{\lambda/\beta}(\mathbf{x}) s_{\lambda/\alpha}(\mathbf{y}).$$

his is true because the Schur polynomials on either side of the above equation (for a skew partition λ and $\lambda' = \mathrm{Cyl}_{k,n}(\lambda)$) describe the ways of filling the same region with letters (since k and $n-k$ are large enough compared to d and the dimensions of α and β). Thus, we have

$$\mathrm{hg}_d \sum_{\mu} s_{\alpha/\mu}(\mathbf{x}) s_{\beta/\mu}(\mathbf{y}) \cdot \sum_{\gamma} s_\gamma(\mathbf{x}) s_\gamma(\mathbf{y}) = \mathrm{hg}_d \sum_{\lambda} s_{\lambda/\beta}(\mathbf{x}) s_{\lambda/\alpha}(\mathbf{y}).$$

This being true for all nonnegative integers d (since we can always find large enough values of k and n), we have proven our proposition. □

The proof style used above can be used to prove other facts about skew tableaux. This means that cylindric tableaux have the potential to be very useful to regular tableau theory.

8. A Note on Knuth Equivalence for Cylindric Tableaux

8.1. Words and Knuth Equivalence

This subsection introduces the concepts of tableau words, Knuth transformations, and Knuth equivalence. These concepts are well-established in tableau theory, but are described here in order to establish conventions. These conventions are the ones used in William Fulton's *Young Tableaux* [Ful, §2.1]. *In this subsection, tableaux will refer to regular (non-cylindric) tableaux.*

Definition 8.1. The *word* of a tableau is the sequence of letters (entries) in the tableau, reading left to right across the tableau's rows and then from bottom to top.

For example, the following (skew) tableau's word is 3346354.

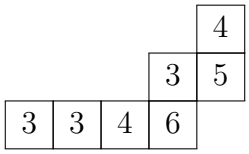

Definition 8.2. Let yzx be a sequence of three consecutive letters of a word w such that $x < y \leq z$. Then a transformation of type K' takes w into the word w' obtained by replacing the three letters yzx with yxz.

For example, we have $3346354 \xrightarrow{K'} 3343654$. Note that K' may not be able to be applied to a word, or might be applicable to a word in more than one way.

Definition 8.3. Let xzy be a sequence of three consecutive letters of a word w such that $x \leq y < z$. Then a transformation of type K'' takes w into the word w' obtained by replacing the three letters xzy with zxy.

For example, we have $3346354 \xrightarrow{K''} 3346534$. Just like K', the transformation K'' may not be able to be applied to a word, or might be applicable to a word in more than one way.

Definition 8.4. An *elementary Knuth transformation* is a transformation that is K', K'', or the inverse transformation of K' or K''. Two words are *Knuth equivalent* if one can be obtained from another through a series of elementary Knuth transformations. Two tableaux are Knuth equivalent if their words are Knuth equivalent.

Knuth equivalence is an equivalence relation; thus, for any two words w and v, if w is equivalent to v, then v is equivalent to w. The significance of Knuth equivalence is that there is a unique semistandard tableau that is Knuth equivalent to a given skew semistandard tableau, and that tableau is that skew tableau's rectification [Ful, §2.1, Corollary 2]. (See [Ful, §1.2] for an explanation of tableau rectification.) In general, Knuth equivalence is an extraordinarily useful notion in tableau theory.

8.2. Cyclic Knuth Equivalence

The concept of tableau words makes sense for cylindric tableaux as well as regular tableaux.

Definition 8.5. The *word* of a cylindric tableau is the sequence of letters (entries) in the tableau, reading left to right across the tableau's rows and then from row k to row 1.

Due to the cyclic nature of cylindric tableaux, it is desirable that two tableaux be Knuth equivalent if one can be obtained from the other by shifting all cells by the same amount in the same direction. For example, the following two tableaux should be Knuth equivalent:

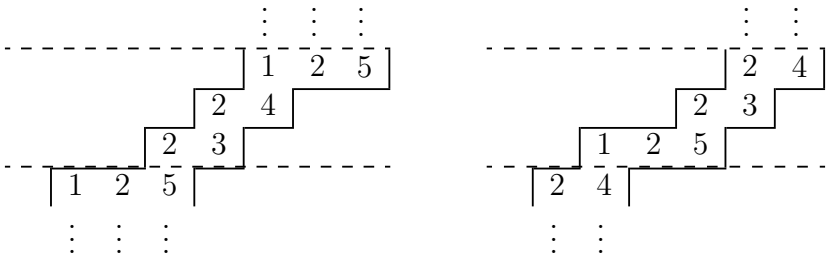

Given our current definition of Knuth equivalence, this is not always the case. For example, the following two tableaux are not Knuth equivalent, as their words (123 and 312, respectively) are not Knuth equivalent.

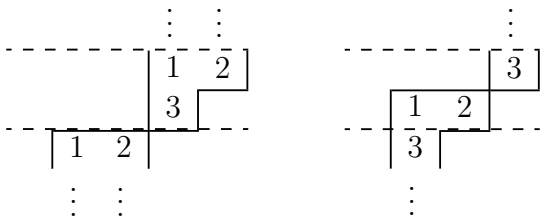

It is natural, then, to define a rotation operation R on a word, which places the rightmost letter of a word to the left of the word. For example, we have $3346354 \xrightarrow{R} 4334635$. By applying R multiple times to a word, we can rotate the letters of the word by any amount.

Definition 8.6. An *elementary cyclic Knuth transformation* is any transformation that is an elementary Knuth transformation or is R. Two words are *cyclic Knuth equivalent* if one can be obtained from another through a series of elementary cyclic Knuth transformation. Two cylindric tableaux are cyclic Knuth equivalent if their words are cyclic Knuth equivalent.

It turns out, however, that cyclic Knuth equivalence is not a useful equivalence relation for cylindric tableau theory, as all tableaux of the same content are cyclic Knuth equivalent. This fact is a consequence of Theorem 5.6.7 in Lothaire's *Algebraic Combinatorics on Words* [Loth, §5.6, Theorem 5.6.7]. Here, we provide a proof of this fact that relies exclusively on basic techniques.

Theorem 8.7. *For any positive integer m, all words of length m that are permutations (i.e. words of length m that consist of the letters 1 through m) are cyclic Knuth equivalent.*

Proof. In this proof, "word" will refer exclusively to permutations. We will also define a word as an equivalence class of words modulo R, since R allows us to rotate any word to any desired position. We will identify each word with its representative that has 1 as its leftmost letter. For any word w, we will say that $w_1 = 1$ and w_i is the letter directly to the right of w_{i-1} for $1 < i \leq m$. In addition, w_1 is directly to the right of w_m.

We will create an algorithm that transforms any word w into $x = 1234\ldots m$. This is sufficient because, for any w and v, if w is equivalent to x and v is equivalent to x, then w is equivalent to v.

Note that the elementary Knuth transformations state that we can switch two consecutive letters if the letter on either side of them is in between the two letters in value.

Definition 8.8. A *catalyst* is a letter that facilitates the swapping of two other letters (i.e. the letter y in definitions 8.2 and 8.3). We say that a catalyst *catalyzes* the swapping of two letters.

Our theorem may be easily verified for $m = 1$, $m = 2$, and $m = 3$. For $m \geq 4$, our algorithm is as follows:

Algorithm 8.9 (Word Transformation Algorithm)**.**

```
    Function WordTrans(word w)          ▷ w must be a permutation of length m.
1:  while w ≠ x do:
2:      i := 1.
3:      while i < m do:
4:          if w_i can switch with w_{i+1} under an elementary Knuth transformation then:
5:              Switch w_i with w_{i+1}.
6:              Exit out of the loop beginning on line 3.
7:          end if.
8:          i := i + 1.
9:      end while.
10: end while.
11: return w.                           ▷ Clearly, x is returned (if this line is reached).
```

An example of this algorithm is shown below, with $w = 159362847$. The two letters that are switched in the following step are bolded. Steps following those in which w_1 switches with w_2 are in red. The reader will gain intuition about the algorithm by following these steps; however, following all of these steps is not necessary for the proof.

15**93**62847 → **15**3962847 → 1**39**628475 → **19**3628475 → 1**36**284759 → **13**2684759 → 1**26847**593 → 1**26**487593 → **16**2487593 → 1**24**875936 → 1**28**475936 → **18**2475936 → 1**24**759368 → 1**27**459368 → **17**2459368 → 1**2459**3687 → 1**24539**687 → 1**24**359687 → **14**2359687 → 1**2359**6874 → 1**23**956874 → 1**29**356874 → **19**2356874 → 1**2368**749 → 1**23**586749 → 1**23**856749 → 1**28**356749 → **18**2356749 → 1**2356**7498 → 1**23564**798 → 1**23**546798 → 1**25**346798 → **15**2346798 → 1**2346**7985 → 1**23469**785 → 1**23496**785 → 1**23**946785 → 1**29**346785 → **19**2346785 → 1**2346**7859 → 1**23467**589 → 1**23465**789 → 1**23**645789 → 1**26**345789 → **16**2345789 → 1**2345**7896 → 1**23457**869 → 1**23457**689 → 1**23475**689 → 1**237**45689 → 1**27**345689 → **17**2345689 → 1**2345**6897 → 1**23456**879 → 1**23458**679 → 1**23485**679 → 1**23**845679 → 1**28**345679 → **18**2345679 → 1**2345**6798 → 1**23456**978 → 1**23459**678 → 1**23495**678 → 1**23**945678 → 1**29**345678 → **19**2345678 → 123456789

Definition 8.10. For any word w, define $M(w)$ to be word $w_1 w_2 \dots w_j$ such that $w_1 < w_2 < \cdots < w_j$, but $w_j > w_{j+1}$ (recall that $w_1 = 1 < w_2$). We will call j the *increase length* of w. (If w has length m and $w_1 < w_2 < \cdots < w_m$ (in other words, $w = x$), then $M(w) = w_1 w_2 \dots w_m$ and the increase length is m.)

With this concept, we can illustrate exactly when the "If" clause of the above algorithm is entered. It is possible that w_1 can be switched with w_2, with w_n as a catalyst. If not, let j be the increase length of w (we assume that $j \neq m$, because in that case we are done). We know that $w_j > w_{j-1}$ and that $w_j > w_{j+1}$, so, out of w_{j-1}, w_j, and w_{j+1}, either w_{j-1} or w_{j+1} is the middle term in value. Thus, w_j can be switched with the smaller of its two neighbors. The switch cannot occur earlier, because two consecutive terms in the middle of w's increasing sequence do not have a term in between them in value to either side.

We conclude that the "If" clause is entered in every pass through the loop beginning on line 1, because $M(w)$ ends with a letter that is not w_m (since we are taking $w \neq x$).

One might observe from the example that the location of the switch always shifts to the left, unless w_1 switched with w_2 in the previous step. This is indeed the case.

Lemma 8.11. *The location of the switch made by the algorithm always shifts to the left, unless w_1 switched with w_2 in the previous step.*

Proof. Let w be a word that is obtained during the execution of Algorithm 8.9 (but not necessarily the word that is taken as input). We will consider two cases.

Case 1: $w = 1cb\dots$, where b and c just switched (that is, w_2 and w_3 just switched values). In this case, c cannot be between 1 and b in value, because then 1 and b would have switched in the previous step. c cannot be less than 1. Thus, $1 < b < c$, which means that b will catalyze the switch between 1 and c.

Case 2: $w = \dots abdc\dots$, where c and d just switched (that is, w_l and w_{l+1} just switched for some $l > 2$). In this case, d cannot be between b and c, because then b and c would have switched in the previous step. We also know that $a < b < c$, because at least one element of the increasing sequence beginning with w_1 is present in every switch (which implies that the increasing sequence continues up to, if not beyond, c). Thus, we have that either $a < b < c < d$, $a < d < b < c$, or $d < a < b < c$. In the first and third cases, b switches with d (which means that the location of the switch shifts one place to the left). In

the second case, a switches with b (which means that the location of the switch shifts two places to the left).

Thus, the location of a switch is always to the left of the location of the previous switch, unless the previous switch was between w_1 and w_2. □

Remark 8.12. *Note that the switch in question always shifts either* 1 *or* 2 *places to the left (if the previous switch was not between w_1 and w_2) as the algorithm proceeds; this fact will be used later.*

From the fact that the switch always shifts to the left, it follows that, from any word, a switch between w_1 and w_2 will eventually happen (unless we arrive at x beforehand). This means that, in performing the algorithm, an infinite loop in which w_1 never switches with w_2 will never be reached. Thus, if there exists an infinite loop in any possible execution of the algorithm, such a loop must contain a word that was obtained by switching w_1 with w_2.

Definition 8.13. A *critical word* is a word obtained by switching w_1 with w_2. Note that no particular word is critical or non-critical; this adjective describes how the word was obtained within a particular execution of Algorithm 8.9.

In the example above, the critical words were written in red. We will examine these critical words more closely.

$\cdots \to$ 139628475 $\to \cdots \to$ 136284759 $\to \cdots \to$ 126847593 $\to \cdots \to$ 124875936 $\to \cdots \to$ 124759368 $\to \cdots \to$ 124593687 $\to \cdots \to$ 123596874 $\to \cdots \to$ 123568749 $\to \cdots \to$ 123567498 $\to \cdots \to$ 123467985 $\to \cdots \to$ 123467859 $\to \cdots \to$ 123457896 $\to \cdots \to$ 123456897 $\to \cdots \to$ 123456798 $\to \cdots \to$ 123456789

For any $1 \leq l \leq m$, let $P_w(l)$ be the integer t such that $w_t = l$. To each word w of length m we can assign a base-$m+1$ number $N(w)$ whose base-$m+1$ representation is $P_w(1)P_w(2)\dots P_w(m)$. For example, $N(159362847) = 164825973$. We will write down the numbers associated with the critical words in our example.

152794863 $\to$ 142683759 $\to$ 129573648 $\to$ 128369547 $\to$ 127358496 $\to$ 126347985 $\to$ 123946875 $\to$ 123845769 $\to$ 123745698 $\to$ 123495687 $\to$ 123485679 $\to$ 123459678 $\to$ 123456978 $\to$ 123456798 $\to$ 123456789

In our example, these numbers strictly decrease. It suffices to prove that they strictly decrease, because then a loop cannot be entered, and the smallest possible value of $N(w)$ for any word w, $123\dots n$, will eventually be reached. It is indeed the case that these numbers are strictly decreasing.

Lemma 8.14. *Let w be a word and let v be any critical word reached by following the algorithm starting with w, except for $123\dots n$. Let v' be the next critical word obtained by the algorithm, assuming that such a word exists.*[11] *Then $N(v') < N(v)$.*

Proof. Let l be the number such that $v_l \neq l$, but for all i such that $1 \leq i < l$, $v_i = i$. Then $N(v) = 123\dots(l-1)\dots$. We have identified the position of the switch that happens to

[11] If v' does not exist, then $123\dots n$ is reached, and so we are done.

v as either between v_1 and v_2, catalyzed by v_m, or between the local maximum[12] with the smallest position number and its smallest adjacent entry.

Case 1: v_1 switches with v_2 with catalyst v_m. Since $v_1 = 1$, v_2 cannot be 2, as nothing can catalyze a switch between 1 and 2. Thus, $N(v)$'s second digit, $P_v(2)$, is some number that does not equal 1 or 2. When v_1 switches with v_2, the position of 2 in the new word (call it v') is smaller by 1 (since $P_{v'}(v_2) = m$, and all other position numbers have decreased by 1). Hence, in this case, $N(v') < N(v)$, as desired.

Case 2: The switch that happens to v is between the local maximum with the smallest position number in v and the smallest entry adjacent to it. The first local maximum cannot be l, as $v_l > l$ and v_l is part of $M(v)$. Furthermore, since l is smaller than some entry to its left, the first local maximum must be to the left of l. This means that the switch that the algorithm performs on v decreases $P_v(l)$ by 1 or keeps it the same. By our discussion above, the switch that happens to the word obtained from v will be to the left of the switch that happens to v, so $P_v(l)$ will either decrease by 1 or stay the same. This will continue until some letter switches with 1, when $P_v(l)$ will decrease by 1 (since the letter occupying the second position now occupies the m'th position). Note that the letter that switches with 1 cannot be l, since, for this to happen, l must have switched with $l-1$ at some point, which is impossible (as there are no letters that could catalyze this switch).

Let v' be the next critical word obtained from v. We have shown that l will have a smaller position number in the next critical word; that is, $P_{v'}(l) < P_v(l)$. Now we will prove that for all $1 \leq i < l$, $P_v(i) = P_{v'}(i)$.

By Remark 8.12, in the process in which v is transforms into v', the location of the switch shifts to the left by 1 or 2 places every time. Combined with the fact that the switch that happens to v is to the right of v_{l-1} (or includes v_{l-1}), we conclude that there is a switch that switches v_{l-1} with another letter during the process that transforms v into v'. Consider the first such switch and let u be the word to which this switch occurs. Note that for all $1 \leq i < l$, $u_i = v_i = i$. Since the switch that happens to u involves $u_{l-1} = l-1$, the switch cannot be between u_{l-1} and u_{l-2} (since, again, two letters whose value differs by 1 cannot be switched). Hence, the switch occurs between u_{l-1} and u_l. Let u' be the word that is obtained after this switch. Then $u' = 123\ldots(l-2)u_l(l-1)$. Since $l-2 < l-1 < u_l$, u_l switches with $l-2$. If the new word is called u'' (and u'' is not a critical word), we have $u'' = 123\ldots(l-3)u_l(l-2)$, so u_l switches with $l-3$. This continues until u_l switches with 1. When this happens, v' is obtained, and by this analysis we know that $P_{v'}(i) = i$ for $1 \leq i < l$.

Thus, we know that $N(v) = 123\ldots(l-1)P_v(l)\ldots$ and $N(v') = 123\ldots(l-1)P_{v'}(l)\ldots$. As proved earlier, $P_{v'}(l) < P_v(l)$. Thus, $N(v') < N(v)$ in this case as well, as desired. □

We have constructed an algorithm that transforms words into other words. We have shown that this algorithm does not get stuck in a loop where no change is made. We have shown that this algorithm does not get stuck in a loop devoid of critical words. We have shown, through a decreasing monovariant, that no critical word can be repeated. Since the word with the smallest monovariant is $123\ldots m$, we have shown that this word is obtained through our algorithm. Since the Knuth transformations are reversible, if one can get from any word p to another word q, one can get from q to p through these transformations. Thus, for any two words w and v, we have shown how to get from w to $123\ldots m$ and from $123\ldots m$ to v,

[12]By this we mean a letter that is greater than the two letters to either side of it.

thus constructing a method to get from w to v. As this is applicable to all words, we have successfully shown that all permutations are cyclic Knuth equivalent. □

We now extend our theorem, which applies only to permutations, to all words, and thus show that any two cylindric tableaux with the same content are cyclic Knuth equivalent.

Corollary 8.15. *For any two words w and v such that v is a permutation of w, w and v are cyclic Knuth equivalent.*

Proof. For any word w of length m that consists of more than one distinct letter (if all letters are the same, then clearly w is cyclic Knuth equivalent to all of its permutations), let $s(w)$ be the smallest letter of w. Construct the word w' as follows: let t be the integer such that:

- If the last letter of w is not $s(w)$, then t is such that the t'th letter of w (from the left) is the leftmost instance of $s(w)$ in w.
- If the last letter of w is $s(w)$, then t is the smallest positive integer greater than 1 such that the t'th letter of w is $s(w)$, but the $(t-1)$'th letter of w is not $s(w)$.

Then, for all $1 \leq i \leq m$, let the i'th letter of w' be the i'th letter of w, plus $\frac{(i-t) \pmod m}{m}$. Next, construct the permutation $p(w)$ of length m such that, for all $1 \leq i, j \leq m$, if the i'th letter of w' is smaller than the j'th letter of w', then the i'th letter of $p(w)$ is smaller than the j'th letter of $p(w)$. For instance, if $w = 43242$, then $s(w) = 2$, $t = 3$, $w' = (4.6)(3.8)(2)(4.2)(2.4)$, and $p(w) = 53142$.

We know that for any word w of length m, $p(w)$ is transformed by Algorithm 8.9 into $x = 12 \ldots m$. Clearly, for any two distinct words w and v that are permutations of each other, $p(w) \neq p(v)$. We show $p(w)$ and $p(v)$ to be equivalent by applying Algorithm 8.9 to each one, transforming them into x. Thus, if we can show that for each switch that occurs during the execution of Algorithm 8.9 on $p(w)$, the corresponding switch of letters in w is permissible under the cyclic Knuth transformations, then it will follow that w and v are cyclic Knuth equivalent. We will now show this.

For $1 \leq i \leq m$, define w_i to be the letter corresponding to $p(w)_i$. (In our example above, the first 2 is w_1.) For any $1 \leq i, j \leq m$, if $p(w)_i < p(w)_j$, then $w_i \leq w_j$. Thus, any restrictions on switches in w that do not exist in $p(w)$ arise only when two letters that are equal in w are part of the switch (either being switched or acting as a catalyst).

The restrictions on switches involving equal letters set by definitions 8.2 and 8.3 can be summarized as follows:

(1) Equal letters cannot be switched;

(2) A letter acting as a catalyst from the left cannot be equal to the smaller of the two letters being switched; and

(3) A letter acting as a catalyst from the right cannot be equal to the larger of the two letters being switched.

Suppose that all switches performed on w corresponding to switches done by Algorithm 8.9 on $p(w)$ have been legitimate up through the formation of a word u (u may equal w).

(We think of w and $p(w)$ as words that change value over time, as opposed to u, which is a particular word.) We will show (assuming that Algorithm 8.9 has not yet terminated) that the switch that takes u to a new word u', analogous to the corresponding switch done by Algorithm 8.9 with input $p(w)$, is also legitimate.

First, we note that, since all switches so far have been legitimate, it is the case that, for all j, the j's in w have stayed in the same order relative to each other. This means that the letters of $p(w)$ corresponding to the j's in w have also stayed in the same order relative to each other. In the original $p(w)$, this order is increasing (in our example above, the letters of $p(w)$ corresponding to the 4's in w are $p(w)_2 = 4$ and $p(w)_4 = 5$). This means that no two of these letters can switch (because, if they are adjacent, they differ by 1, which means that no letter can catalyze their switch). Consequently, condition (1) is satisfied for the switch that takes u to u'.

Suppose that condition (2) is violated by the switch that would take u to u'. Due to the increasing nature of the letters of $p(w)$ corresponding to the j's in w (for any j) and the fact that $p(w)_1$ never switches with $p(w)_m$ during the execution of Algorithm 8.9, the only possibility of this is if $p(w)_m$ catalyzes the switch between $p(w)_1$ and $p(w)_2$. In order for this scenario to be a violation of condition (2), u_m must equal $s(w)$ (note that although w changes, $s(w)$ is constant). We now show that this cannot be the case.

Suppose that $u_m = s(w)$. For the original value of w, $w_m \neq s(w)$ (because t is originally chosen such that the $(t-1)$'th letter of w — that is, w_m — is not $s(w)$). Thus, w_m becomes $s(w)$ after some switch that happens to a value of w that precedes u. In order for this to happen, one of two things must happen: $p(w)_1$ must switch with $p(w)_2$, with $p(w)_2$ corresponding to $s(w)$ in w, — this is impossible, as previously discussed — or $p(w)_m$ must switch with $p(w)_{m-1}$, with $p(w)_{m-1}$ corresponding to $s(w)$ in w. This cannot be the case, because neither $p(w)_{m-1}$ nor $p(w)_m$ can be the first local maximum in $p(w)$ (as we discussed in our proof of Theorem 8.7, every switch not between $p(w)_1$ and $p(w)_2$ involves the first local maximum). Thus, $u_m \neq s(w)$ and condition (2) is not violated.

Suppose that condition (3) is violated by the switch that would take u to u'. Due to the increasing nature of the letters of $p(w)$ corresponding to the j's in w (for any j) and the fact that $p(w)_1$ never switches with $p(w)_m$ during the execution of Algorithm 8.9, the only possibility of this is if $p(w)_1$ catalyzes a switch between $p(w)_{m-1}$ and $p(w)_m$; this is clearly impossible, as $p(w)_1$ cannot act as a catalyst.

Having proven that no condition is violated by the switch that takes u to u', we have completed our induction and have shown that w and v are cyclic Knuth equivalent for all w and v that are permutations of each other. □

It would be very helpful to have an analog to Knuth equivalence for cylindric tableaux; however, cyclic Knuth equivalence is clearly not the desired analog.

References

[Ful] Fulton, W. (1997). *Young tableaux.* Cambridge, United Kingdom: Cambridge University Press.

[GesKra] Gessel, I. M., & Krattenthaler, C. (1997). Cylindric partitions. *Transactions of the American Mathematical Society, 349*(2), 429–479. Retrieved from http://www.ams.org/journals/tran/1997-349-02/S0002-9947-97-01791-1/S0002-9947-97-01791-1.pdf

[GriRei] Grinberg, D., & Reiner, V. (2015, June 1). *Hopf algebras in combinatorics.* Retrieved from http://arxiv.org/pdf/1409.8356v2.pdf[13]

[Knu] Knuth, D. E. (1970). Permutations, matrices, and generalized Young tableaux. *Pacific Journal of Mathematics, 34*(3), 709–727. Retrieved from http://projecteuclid.org/download/pdf_1/euclid.pjm/1102971948

[Loth] Lascoux, A., Leclerc, B., & Thibon, J.-Y. (2002). The plactic monoid. In M. Lothaire (Author), *Algebraic combinatorics on words* (pp. 144–172). Cambridge, United Kingdom: Cambridge University Press. Retrieved from http://www-igm.univ-mlv.fr/~berstel/Lothaire/AlgCWContents.html

[McN] McNamara, P. (2006). Cylindric skew Schur functions. *Advances in Mathematics, 205*(1), 275–312. http://dx.doi.org/10.1016/j.aim.2005.07.011

[MorSch] Morse, J., & Schilling, A. (2012). A combinatorial formula for fusion coefficients. *DMTCS Proceedings, AR*, pp. 735–744. Retrieved from http://www.dmtcs.org/pdfpapers/dmAR0165.pdf

[Post] Postnikov, A. (2005). Affine approach to quantum Schubert calculus. *Duke Mathematical Journal, 128*(3), 473–509. Retrieved from http://www-math.mit.edu/~apost/papers/affine_approach.pdf

[SagStan] Sagan, B. E., & Stanley, R. P. (1990). Robinson-Schensted algorithms for skew tableaux. *Journal of Combinatorial Theory, Series A, 55*(2), 161–193. Retrieved from http://dx.doi.org/10.1016/0097-3165(90)90066-6

[13]For solutions to the exercises in this paper, see https://web.archive.org/web/20150606231711/http://www.cip.ifi.lmu.de/~grinberg/algebra/HopfComb-sols.pdf.

Acknowledgments

First and foremost, I would like to thank my mentor, MIT graduate student Darij Grinberg, for introducing me to this topic, answering all of my questions, and helping to proofread my paper. I would also like to thank MIT professors Slava Gerovitch, Pavel Etingof, and Tanya Khovanova for organizing and managing PRIMES, the program that facilitated my research. Finally, I would like to thank MIT professor Alexander Postnikov for suggesting this subject of research.